Past, Present and Future Hammer in Context

Fred Weber

Vice President and CTO, Computation Products Group

Advanced Micro Devices

The Road to Hammer

AMD DEVELOPER symposium 2002

- 286, 386, 486
  - Second source
  - Strong manufacturing, value added engineering
  - Intel ISA, architecture, Design, Infrastructure
  - Some steps on Design – cache sizes and policy
- AMD-K5 and AMD-K6®
  - Independent designs
  - Price/performance leadership
  - Enabled the sub $1000 PC revolution
  - Intel ISA and Infrastructure
  - Some steps on ISA – 3DNow!™ technology
  - Established a strong consumer retail presence and OEM acceptance

The Road to Hammer

AMD DEVELOPER symposium 2002

- AMD Athlon™ processor
  - Performance AND price/performance leadership
    - First time ever that Intel is truly unseated technically
  - Independent infrastructure
    - S2K bus
    - Rich cooperation with third parties – virtual gorilla
      - Chipsets
      - Motherboards
      - Software
  - Consolidation of consumer market
  - Strong inroads to commercial, mobile and server/workstation market
  - >20% market share
  - AMD is here to stay

Hammer

AMD DEVELOPER symposium 2002

- Owners of the x86 roadmap
  - X86-64 Instruction set a clear winner
- Strong systems architecture
  - Solutions for mobile, desktop and server/workstation
    - Glue less MP, HyperTransport™ IO
- Strong virtual gorilla
  - Multiple chipsets and motherboards planned around launch
  - Little question of product acceptance
  - Many server systems in development
  - Tremendous software story
    - Designed for 100% compatibility with existing 32-bit software for x86 instructions
    - Strong 64 bit OS and tools story
    - Tremendous 32 bit app on 64 bit system story
    - Strong interest in 64 bit app porting

The Road Ahead

AMD DEVELOPER symposium 2002

- X86-64 as the volume 64 bit solution
- Ubiquitous 64 bit platforms
- Consumer and Commercial
- Desktop, Mobile, Workstation and Server
- Bigger servers
- Communications/Server convergence

Innovation
U.S. Patents, 1999-2001

AMD DEVELOPER symposium 2002

AMD ranked fifth among US firms during last two years.

Rank	Company	# of Technology Patents in 2001
1	IBM	3,453
2	NEC	1,966
3	CANON	1,877
4	MICRON TECHNOLOGY	1,643
5	SAMSUNG	1,451
6	MATSUSHITA	1,447
7	SONY	1,392
8	HITACHI	1,283
9	MITSUBISHI	1,210
10	FUJITSU	1,208
11	TOSHIBA	1,176
12	LUCENT	1,119
13	GENERAL ELECTRIC	1,112
14	**ADVANCED MICRO DEVICES**	**1,090**
15	HEWLETT-PACKARD	982
16	PHILIPS	882
17	SIEMENS	817
18	**INTEL**	811
19	TEXAS INSTRUMENTS	806
20	MOTOROLA	785
21	XEROX	722
22	EASTMAN KODAK	719
23	BOSCH	705
24	UMC	596
25	HONDA	587

In 2000: 1,055

In 1999: 825

In 2000: 797

In 1999: 735

(Source: IFI Claims)

Source: IFI Claims

An Example – CPU Design Present Situation

AMD DEVELOPER symposium 2002

- AMD-K6® was built entirely on Sparc, PA-RISC and Power machines
- AMD Athlon™ was built 50% on AMD-K6 processors running Linux
  - Few apps. Mostly only ran in-house logic simulators
  - Hammer was built 80% on AMD Athlon processors running Linux
    - Many apps available. Only 64 bit apps conspicuously missing
- Hardware
  - Over 3000 AMD Athlon CPUs doing back-end CAD work in California and Austin
  - Over 1500 AMD Athlon CPUs doing front-end design world-wide
  - Non-AMD processor-based machines are used only for applications which require more memory than x86 is capable of addressing
  - Software
    - Predominantly Linux based
    - Transitioning away from non-x86 based Unix (Solaris, HP-UX, etc.)
    - 64-bit software is run on non-AMD processor-based machines

Future Situation – Next Core Tapeout

AMD DEVELOPER symposium 2002

- We expect to tapeout next generation core using **only** AMD Opteron™ Processors

- Hardware
  - Create a homogenous compute environment
  - Anticipate over 8000 AMD Opteron and AMD Athlon™ CPUs doing back-end CAD work in Sunnyvale and Austin
  - Anticipate over 2000 AMD Opteron and AMD Athlon CPUs doing front-end design world-wide
  - AMD **will not** use any non-AMD 32-bit or 64-bit hardware
- Software
  - 32-bit and 64-bit applications running side by side
  - Large memory applications planned to scale well on AMD Opteron – **4P = 16-32 GB of RAM**

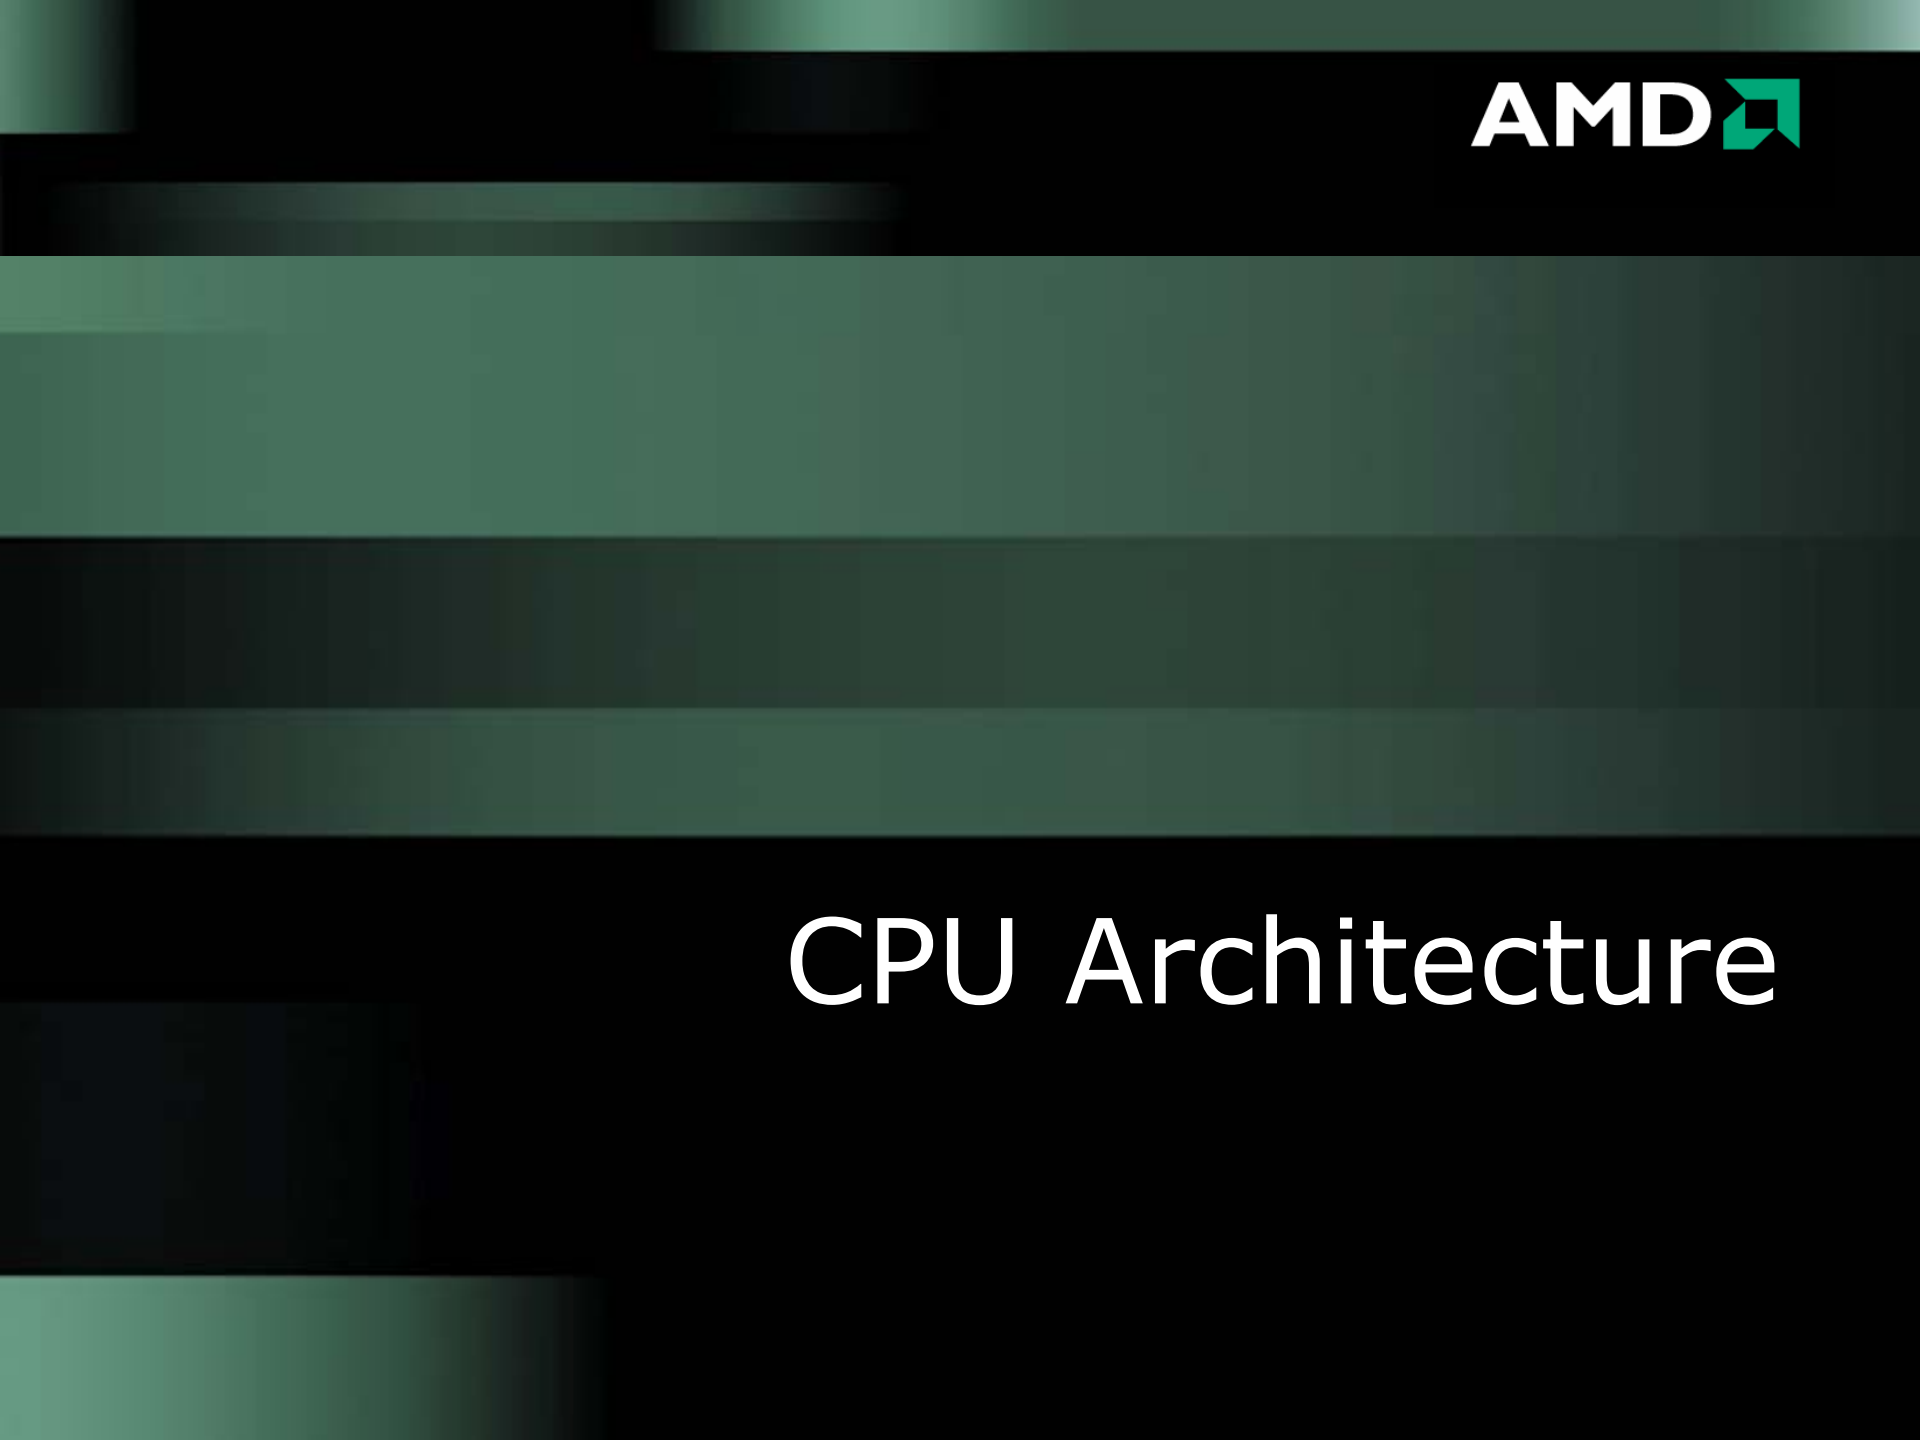
AMD
CPU Architecture

Hammer Architecture Overview

- First x86-64 based processor
- Aggressive out-of-order, 9-issue superscalar processor
- Integrated DDR memory controller
- Leading performance in integer, floating point and multimedia
  - x86-64, x87, MMX™, 3DNow!™, SSE, SSE2

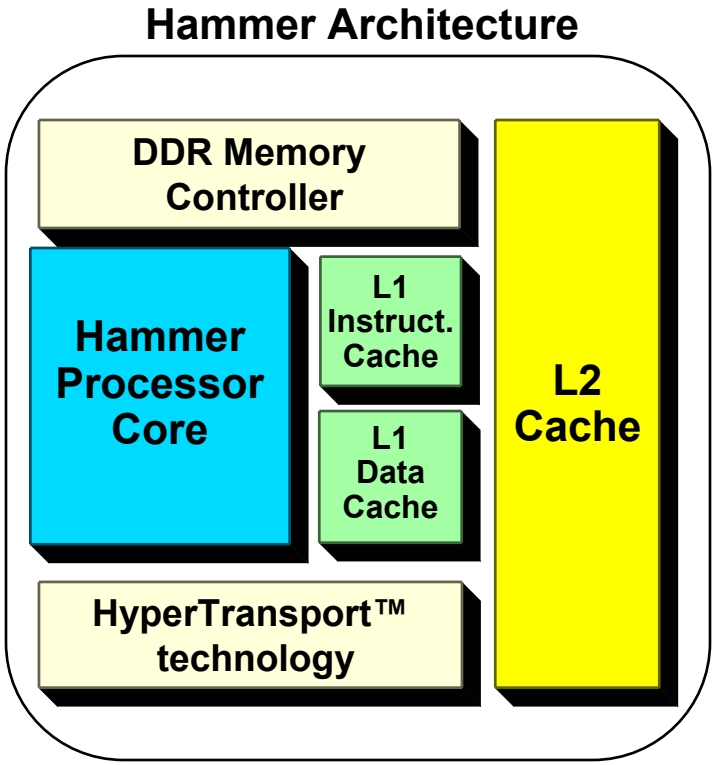

Hammer Core Overview

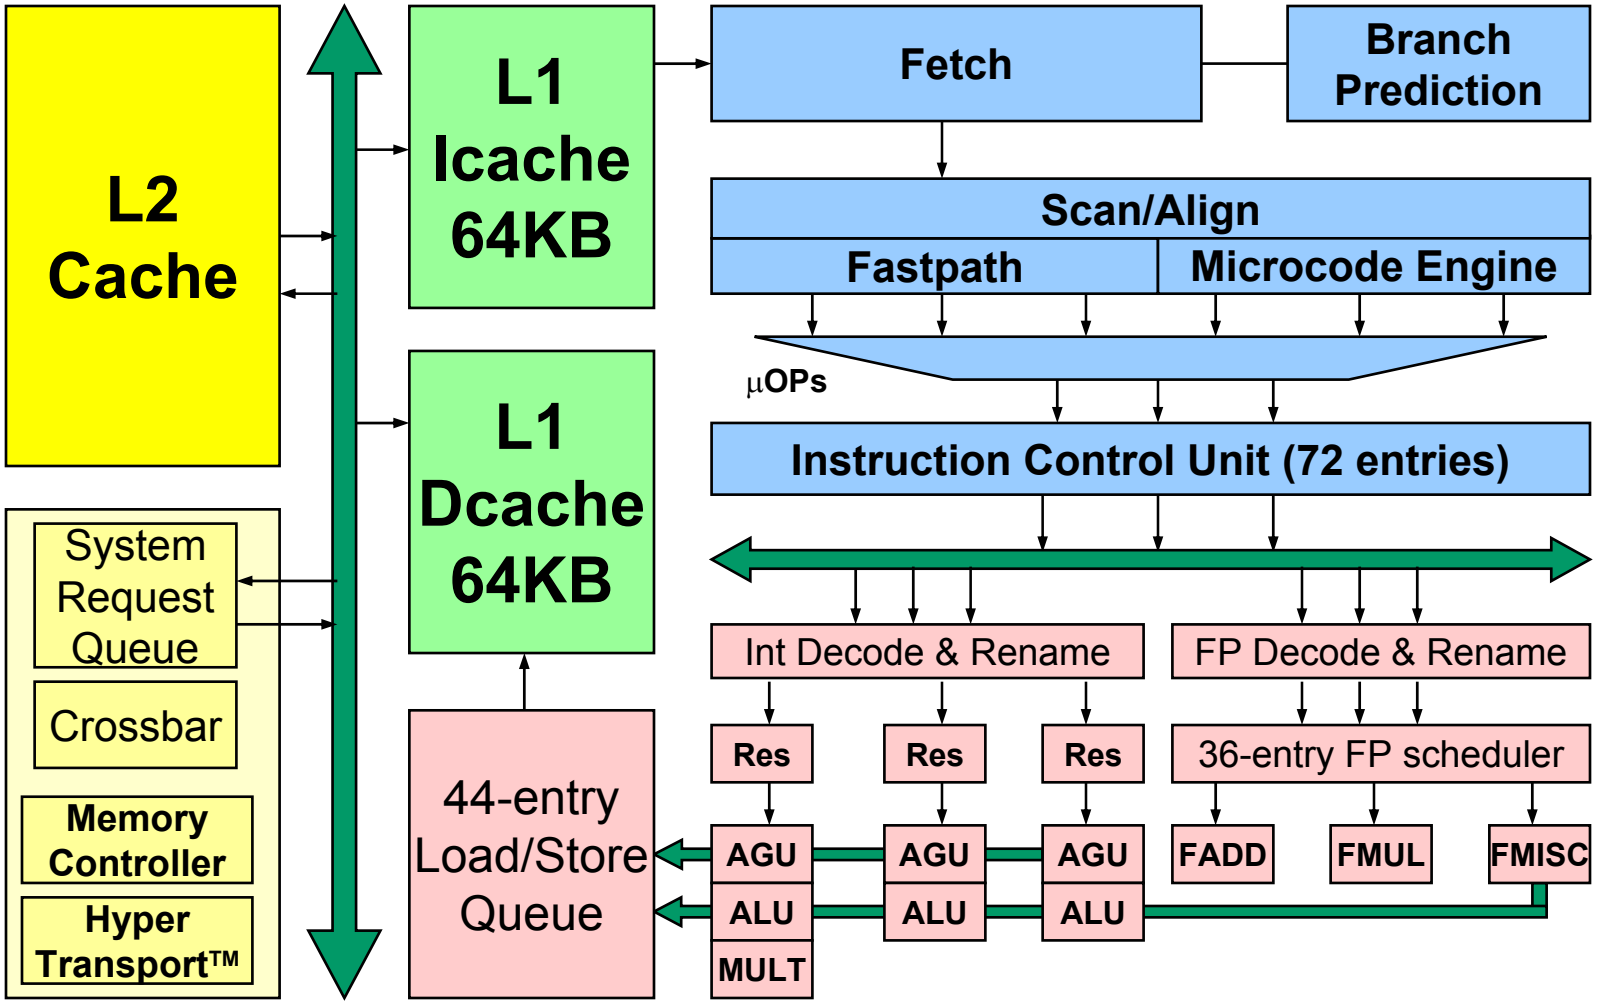

Instruction Fetch

- Supply 16 instruction bytes to the decoder per cycle
- 64KB instruction cache, 2-way set associative
  - Linearly-indexed, physically-tagged, 64-byte block size
  - Prefetch next sequential block on a miss
- 2 sets of instruction cache tags (fetch port, snoop)
- Predecode instruction
  - 1 end bit per-byte
  - Decode some branch types
- Branch prediction

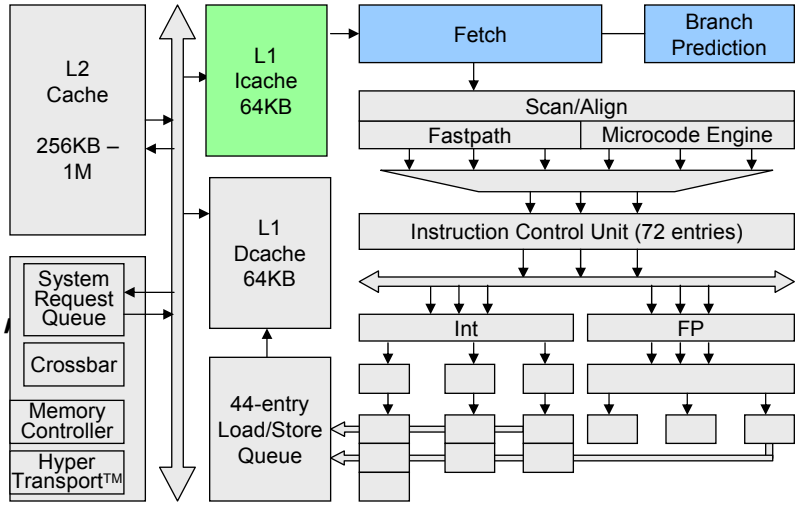

Branch Prediction

- Sequential Fetch

- Predicted Fetch

- Branch Target Address Calculator Fetch

- Mispredicted Fetch

- 5-10% improvement in prediction accuracy vs. AMD Athlon™

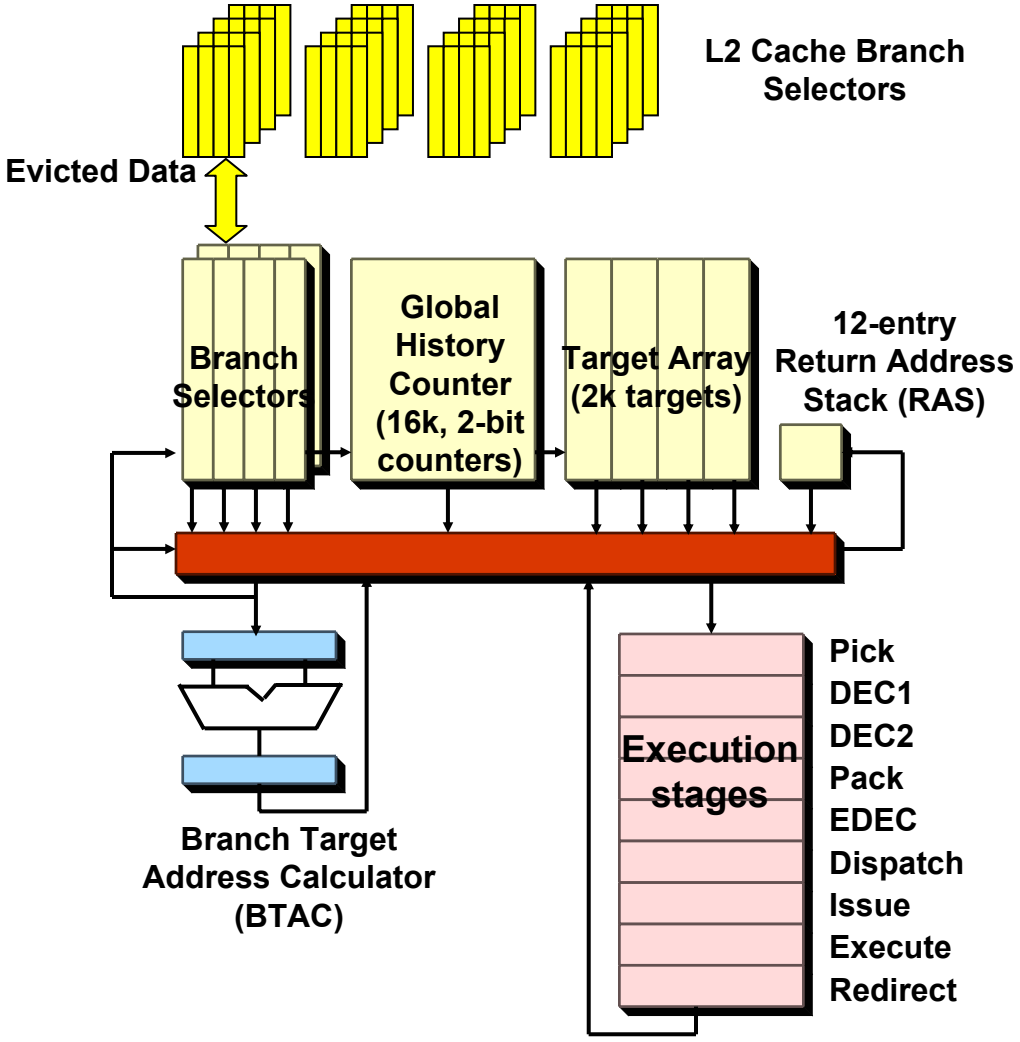

Scan / Align

- Convert x86 instructions to fixed length μOPs
- Dispatch 3 μOPs per cycle to integer/FP schedulers
- Instructions use one of two decoding pipelines
  - Fastpath: instructions decoding to two or fewer μOPs are decoded by hardware, packed into 3 dispatch positions
  - Microcode: x86 instructions decoding to more than two μOPs, calculate ROM entry point, fetch sequence from ROM
- Compared to AMD Athlon™, more instructions use the fastpath
  - Eg: Packed SSE is microcoded in AMD Athlon and fastpath in Hammer
  - Hammer has 8% fewer microcoded instructions for Specint2000
  - Hammer has 28% fewer microcoded instructions for Specfp2000

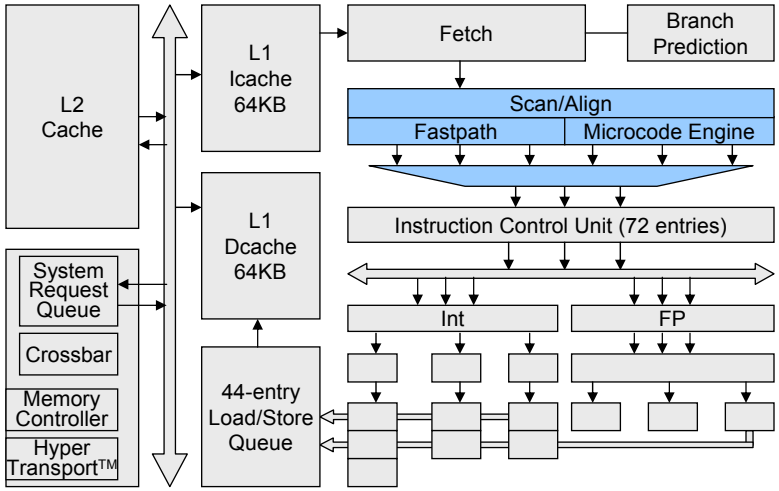

Execution Units

- 3 integer units
- 3 address generation units
- 3 superscalar floating point units
- Integer
  - Full 64-bit data path
  - 3 x 8-entry reservation stations
  - Single cycle 32 and 64-bit add, sub, rotate, shift, logical, etc.
  - 32-bit multiply: 3 cycle latency
  - 64-bit multiply: 5 cycle latency
- Floating point
  - Handles x87, MMX™, 3DNow!™, SSE and SSE2
  - 36-entry scheduler
  - Out-of-order, fully pipelined design

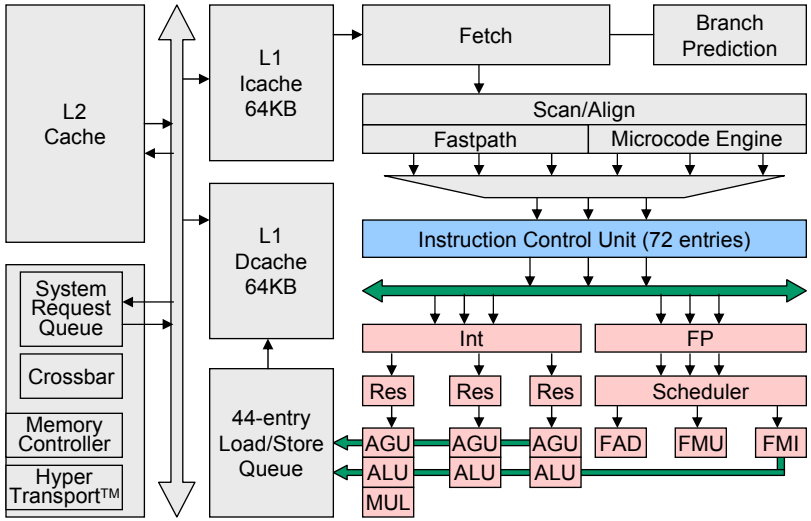

Load/Store and Data Cache

- 64KB data cache
  - 2-way set associative
  - Linearly-indexed, physically-tagged
  - 40-bit physical address
  - 48-bit linear address
  - MOESI coherency
  - 64-byte block size
- Banked and dual ported
  - 2 64-bit reads/writes each cycle to different banks
- 3 sets of data cache tags (port A, port B, snoop)
- Load->use latency is 3 cycles (zero segment base)
  - 1 extra cycle to handle misaligned (quadword boundary) loads
- Data forwarding from stores to dependent loads
- Hardware prefetch

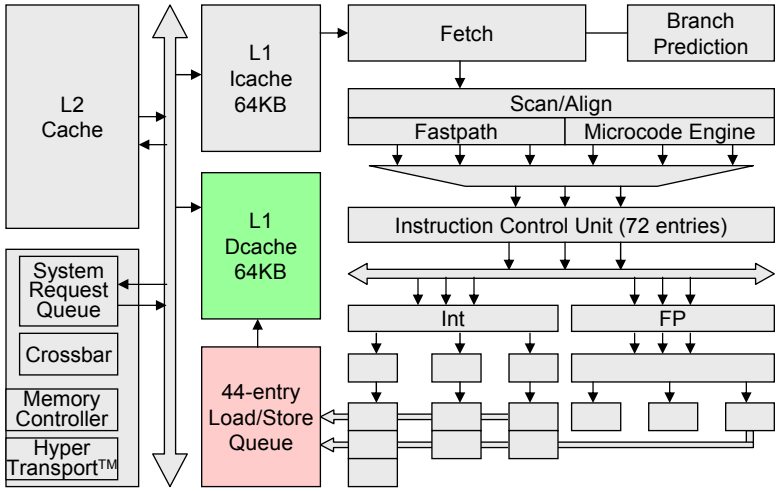

L2 Cache

- Configurable sizes up to 1MB
- 16-way set associative
- L1 and L2 storage is mutually exclusive
- Pseudo-LRU scheme to reduce the number of LRU bits by half
- Stores IC predecode and branch prediction bits
- 10 outstanding miss requests
  - 8 DC
  - 2 IC
- System interface
  - Victim Buffer (8-entry)
  - Snoop Buffer (8-entry)
  - Write Buffer (4-entry)

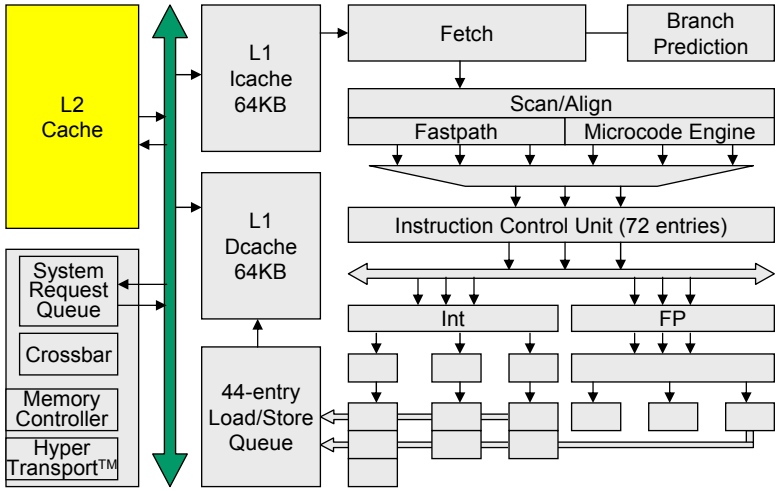

TLB for Large Workloads

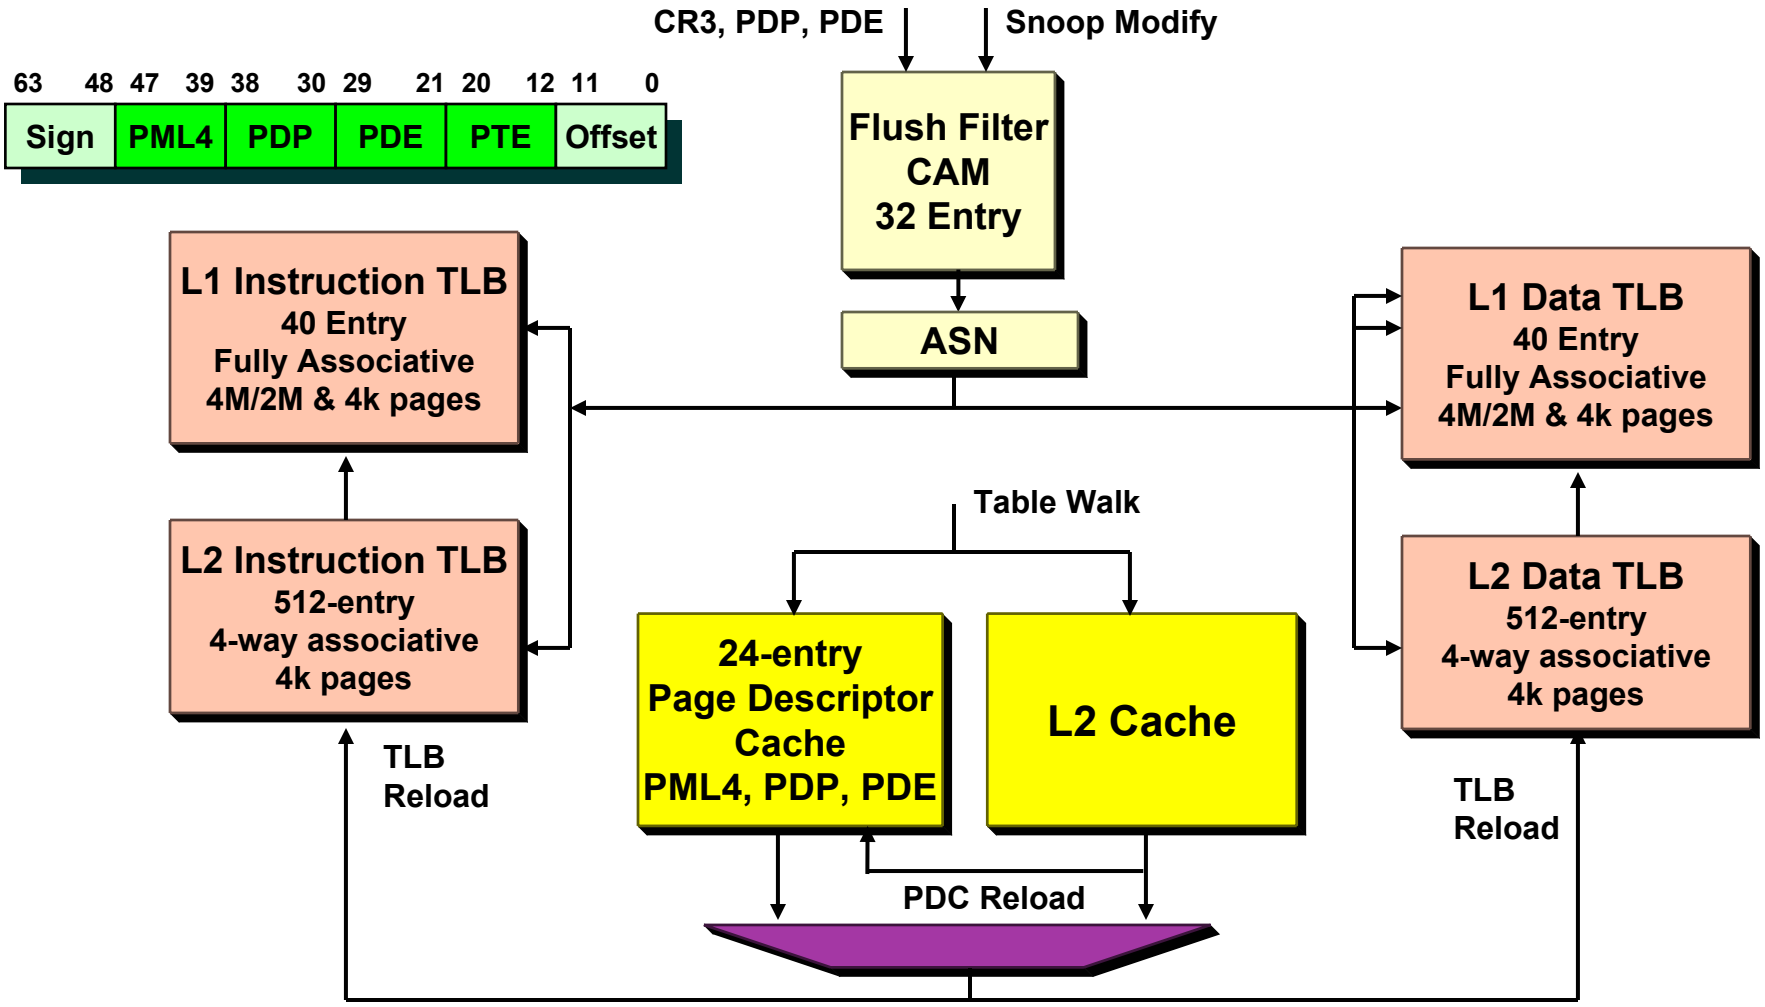

Integrated Memory Controller

AMD DEVELOPER symposium 2002

- Integrated DDR memory controller
  - 8-byte or 16-byte interface
  - Unbuffered or Registered DIMMs
  - 16-byte interface supports direct connection to 8 registered DIMMs and chipkill ECC
  - Significantly reduces memory latency
  - Memory latency improves as CPU and HyperTransport™ link speed improves
  - Performance improves by approximately 20% compared to AMD Athlon™ topology
  - Snoop throughput scales with CPU frequency
- Integrated Northbridge Functionality
  - Processes requests from CPU/IO to DRAM/IO
  - HyperTransport routing
    - peak bandwidth = 6.4GB/s
  - Handles transaction ordering and cache coherence
  - Runs at the same frequency as CPU core

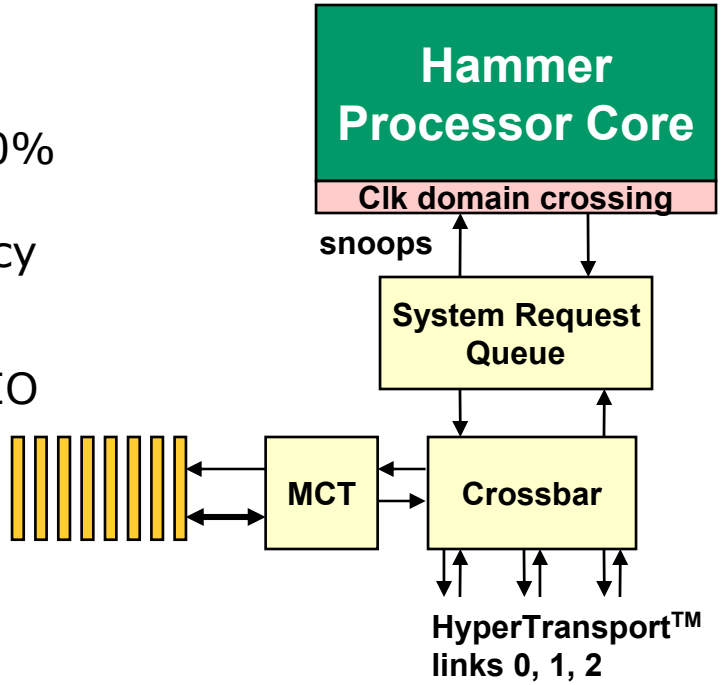

Summary

AMD DEVELOPER symposium 2002

- AMD's Hammer architecture provides a foundation for market-specific solutions:
  - Desktop, mobile, workstation (1-2 way), server (1-8 way)
- AMD's next-generation microprocessor core
  - Leading edge performance for 16-, 32- and 64-bit applications
- Cache subsystem
  - Enhanced TLB structures
  - Improved branch prediction
  - Reliability features
- Integrated DDR memory controller
  - Reduced memory latency
  - 1:1 scaling
- HyperTransport™ technology
  - Fast I/O for chip-to-chip communication
  - Enables glueless MP
- Innovation with compatibility

System Architecture

AMD Opteron™ Processor Architecture

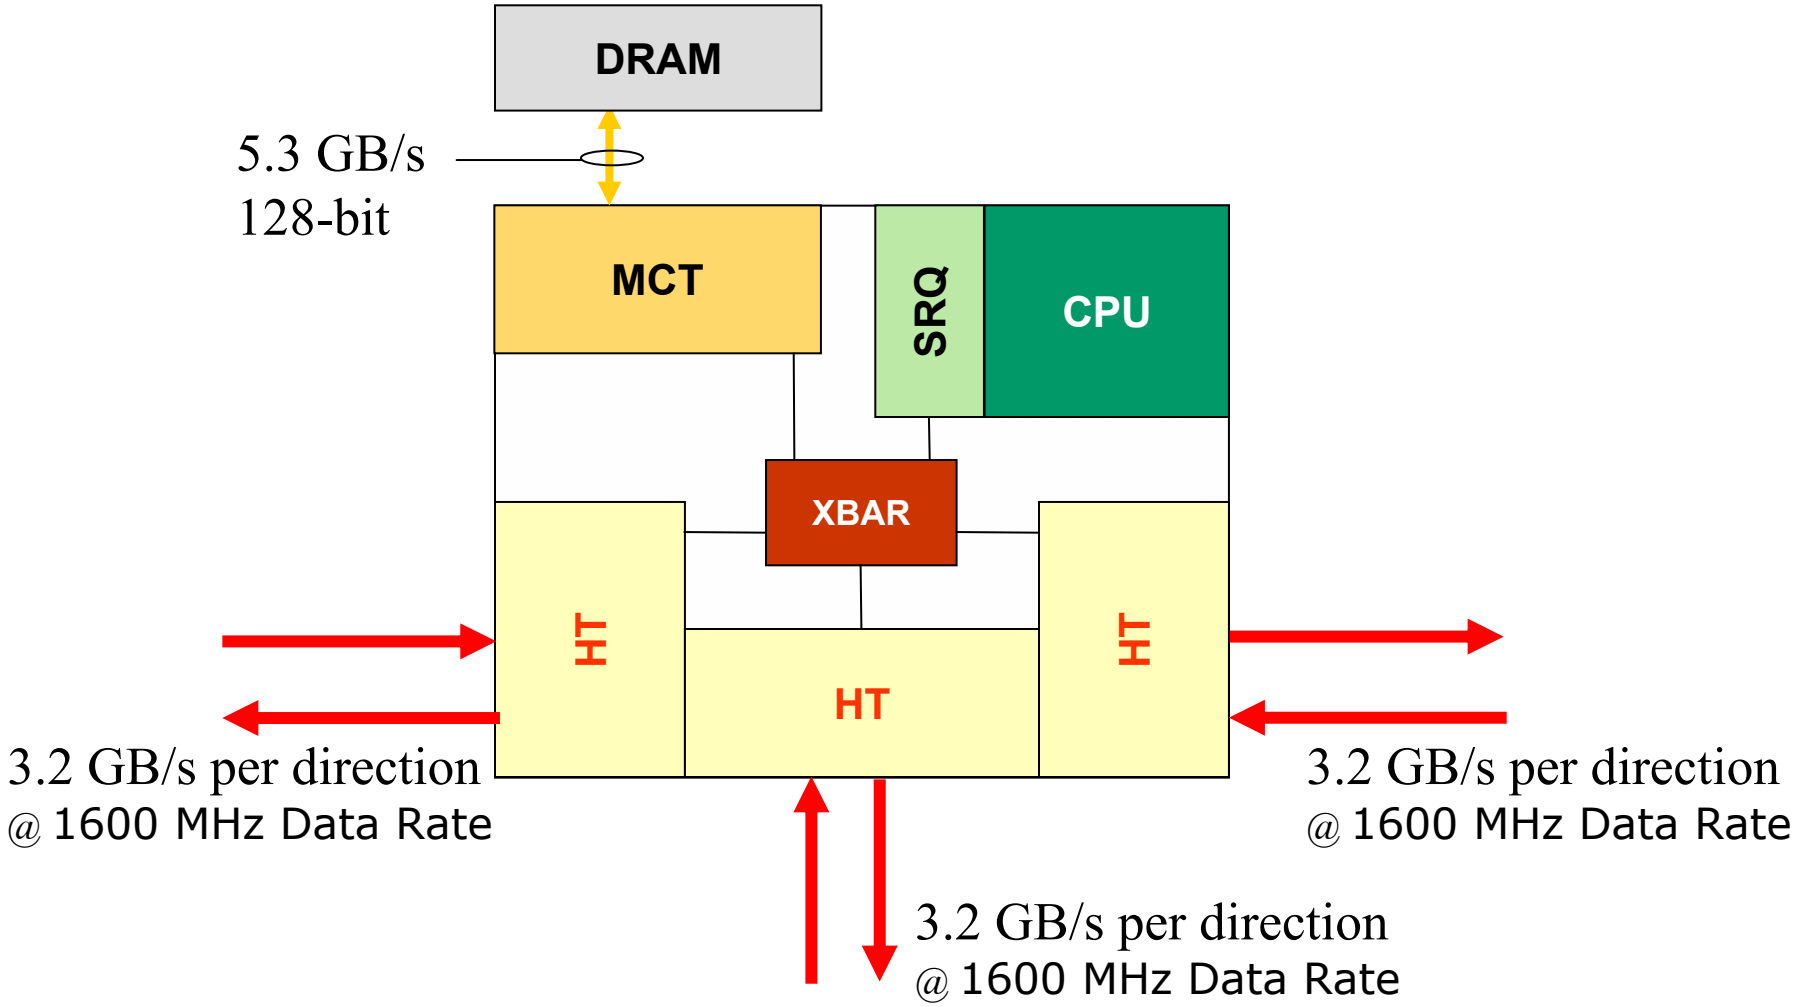

HT = HyperTransport™ technology

Glueless MP System

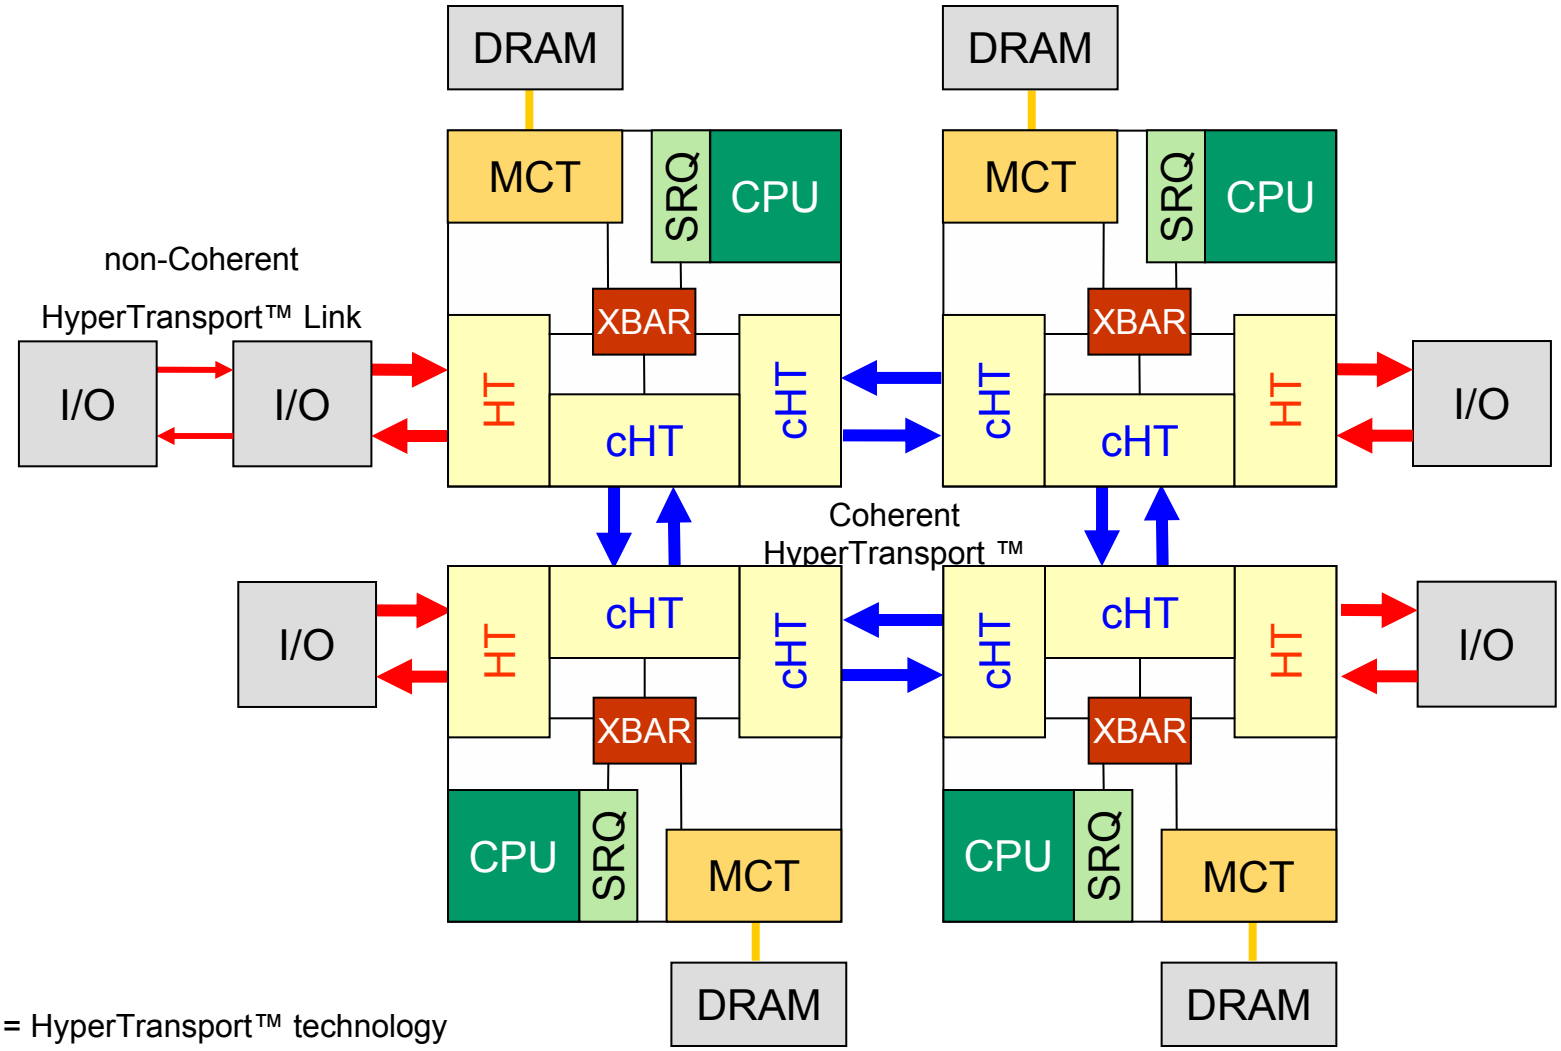

HT = HyperTransport™ technology

MP Architecture

AMD DEVELOPER symposium 2002

- Programming model of memory is effectively SMP
  - Physical address space is flat and fully coherent
  - Far to near memory latency ratio in a 4P system is designed to be <1.4
  - Latency difference between remote and local memory is comparable to the difference between a DRAM page hit and a DRAM page conflict
  - DRAM locations can be contiguous or interleaved
  - No processor affinity or NUMA tuning required
- MP support designed in from the beginning
  - Lower overall chip count results in outstanding system reliability
  - Memory Controller and XBAR operate at the processor frequency
  - Memory subsystem scale with frequency improvements

MP Architecture (contd.)

AMD DEVELOPER symposium 2002

- Integrated Memory Controller
  - 333 MHz 128-bit DRAM interface with up to 8 registered DIMMs
  - High-bandwidth (5.3 GB/s peak) and low-latency memory access
  - Snoop throughput scales with Processor frequency
  - Broadcast cache coherence protocol
    - Avoids serialization delay of directory based systems
    - Snooping the processors caches is overlapped with DRAM access

HyperTransport™ Technology

- Fast I/O for chip-to-chip communication
  - High bandwidth
  - Point-to-point links
  - Split transaction and full duplex
  - Differential Signaling
  - Tunneling capability
- Enables scalable 2-8 processor Cache-Coherent MP systems
  - Glueless MP
- HyperTransport™ Links
  - Up to three 16-bit links (3.2 GB/s per direction)
  - Reduced pin count compared to the typical Bus based systems
  - Compatible with high-volume PC board infrastructure
  - Each can be:
    - cHyperTransport: coherent (Processor-to-Processor) link or,
    - HyperTransport: non-coherent (Processor-to-I/O) link
  - For more info see: http://www.HyperTransport.org/

Local vs. Remote memory access

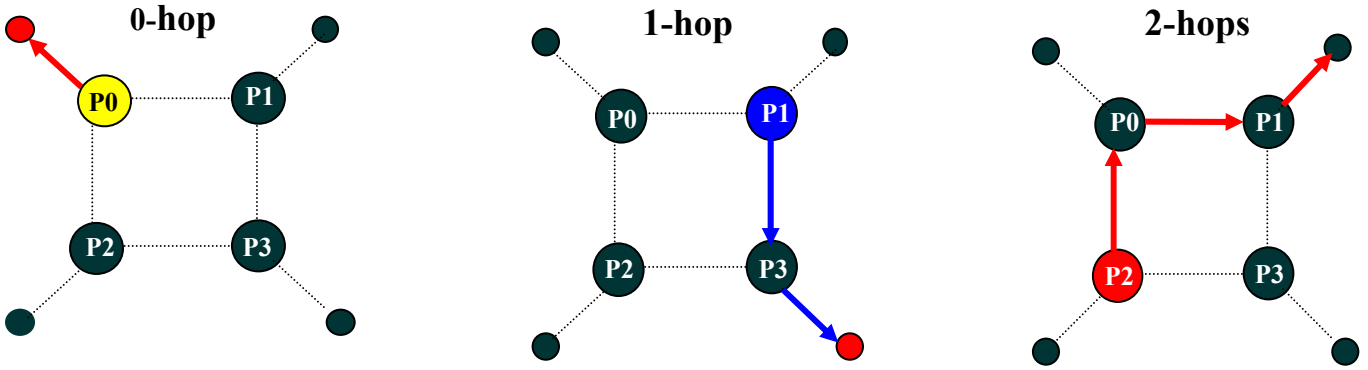

- Local Memory Access (0-hop)
- Remote1 Memory Access (1-hop)
- Remote2 Memory Access (2-hops)

Cache Coherence Protocol Read Transaction Example

AMD DEVELOPER symposium 2002

Step 1

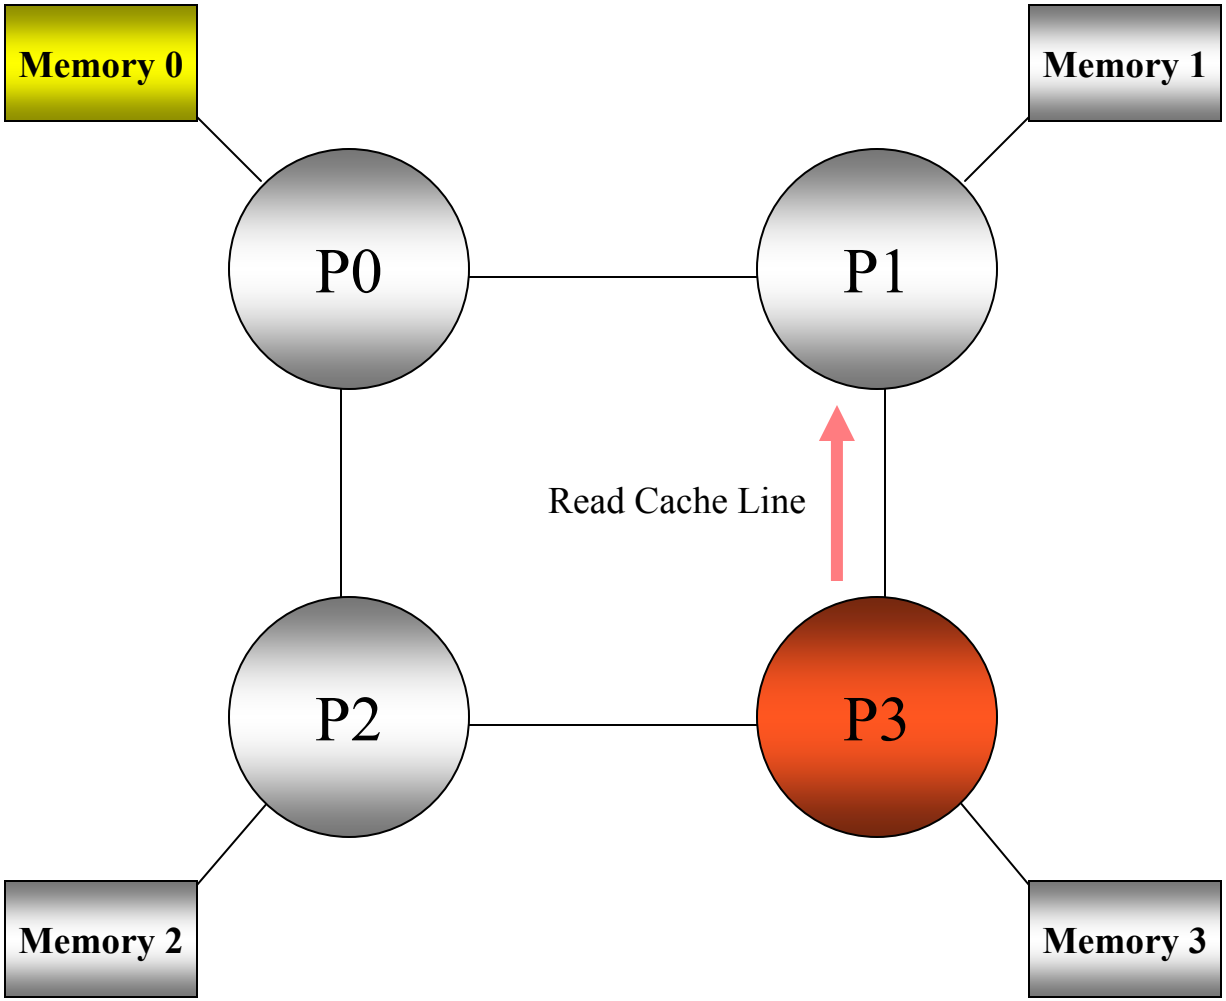

Cache Coherence Protocol Read Transaction Example

Step 2

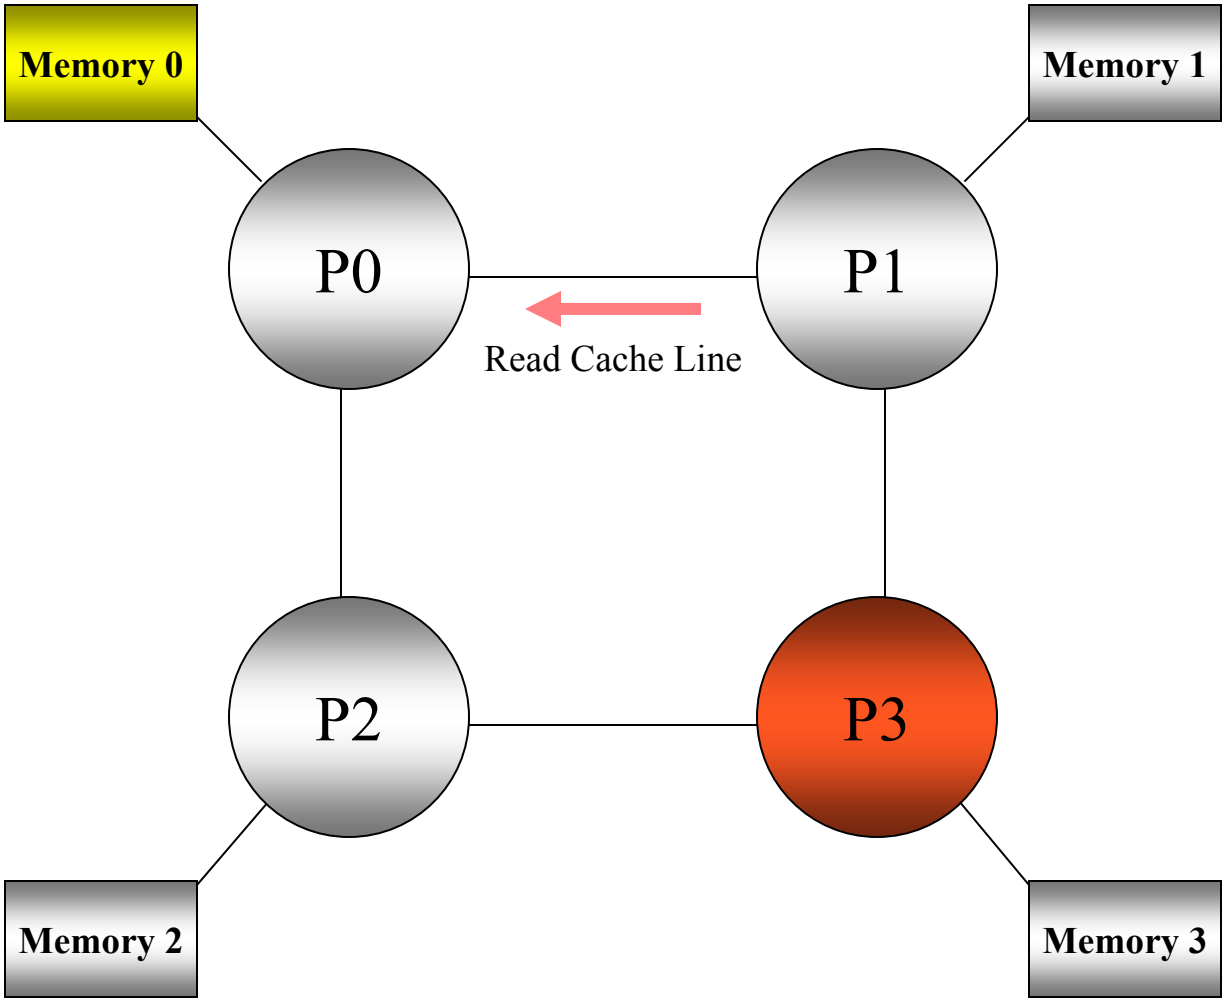

Cache Coherence Protocol Read Transaction Example

Step 3

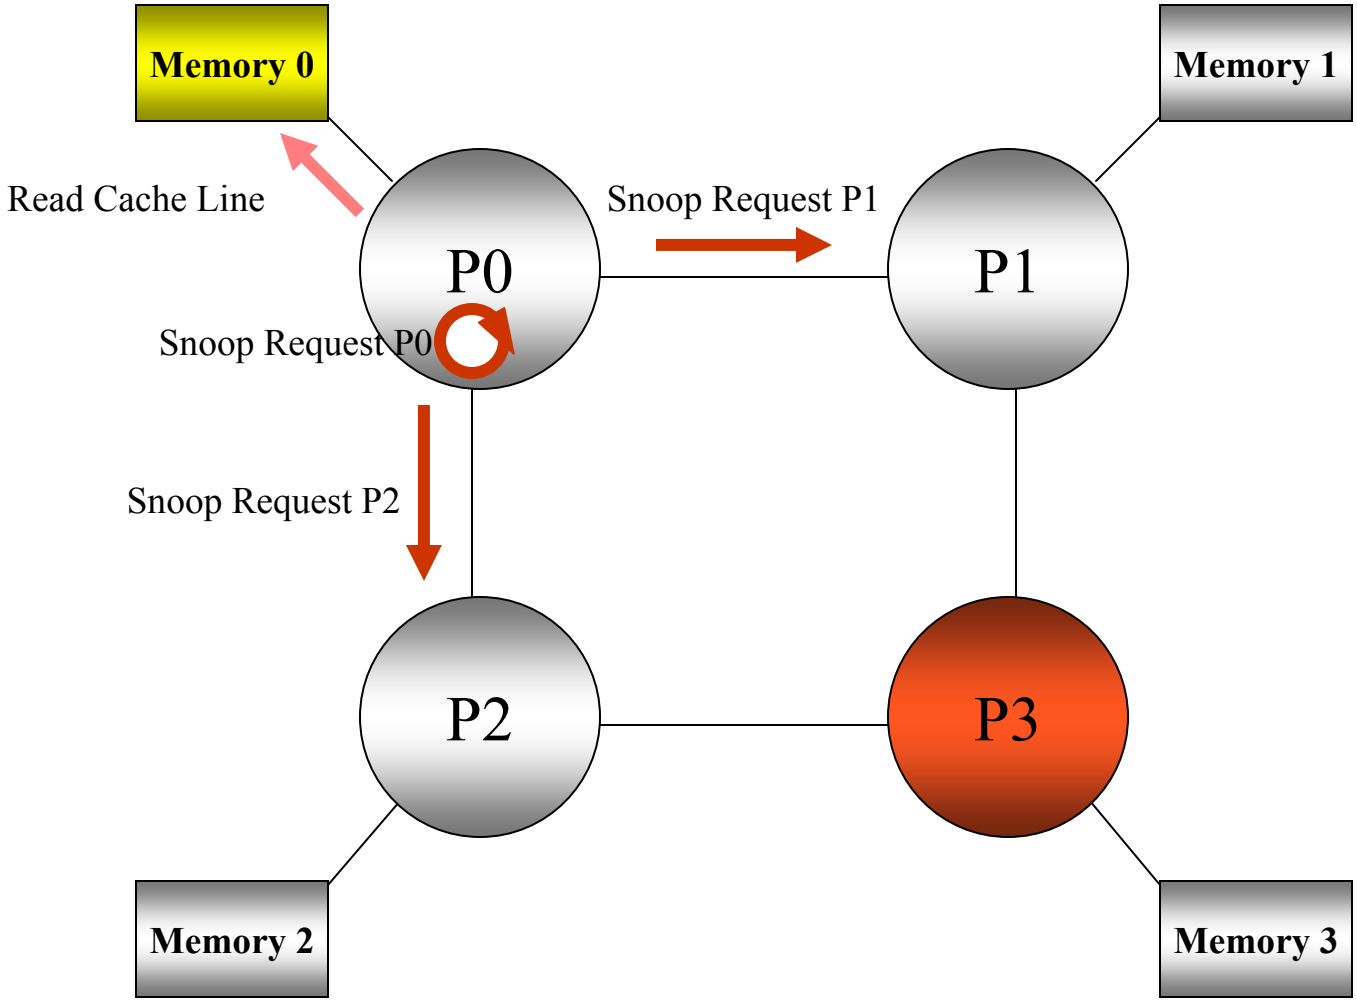

Cache Coherence Protocol Read Transaction Example

Step 4

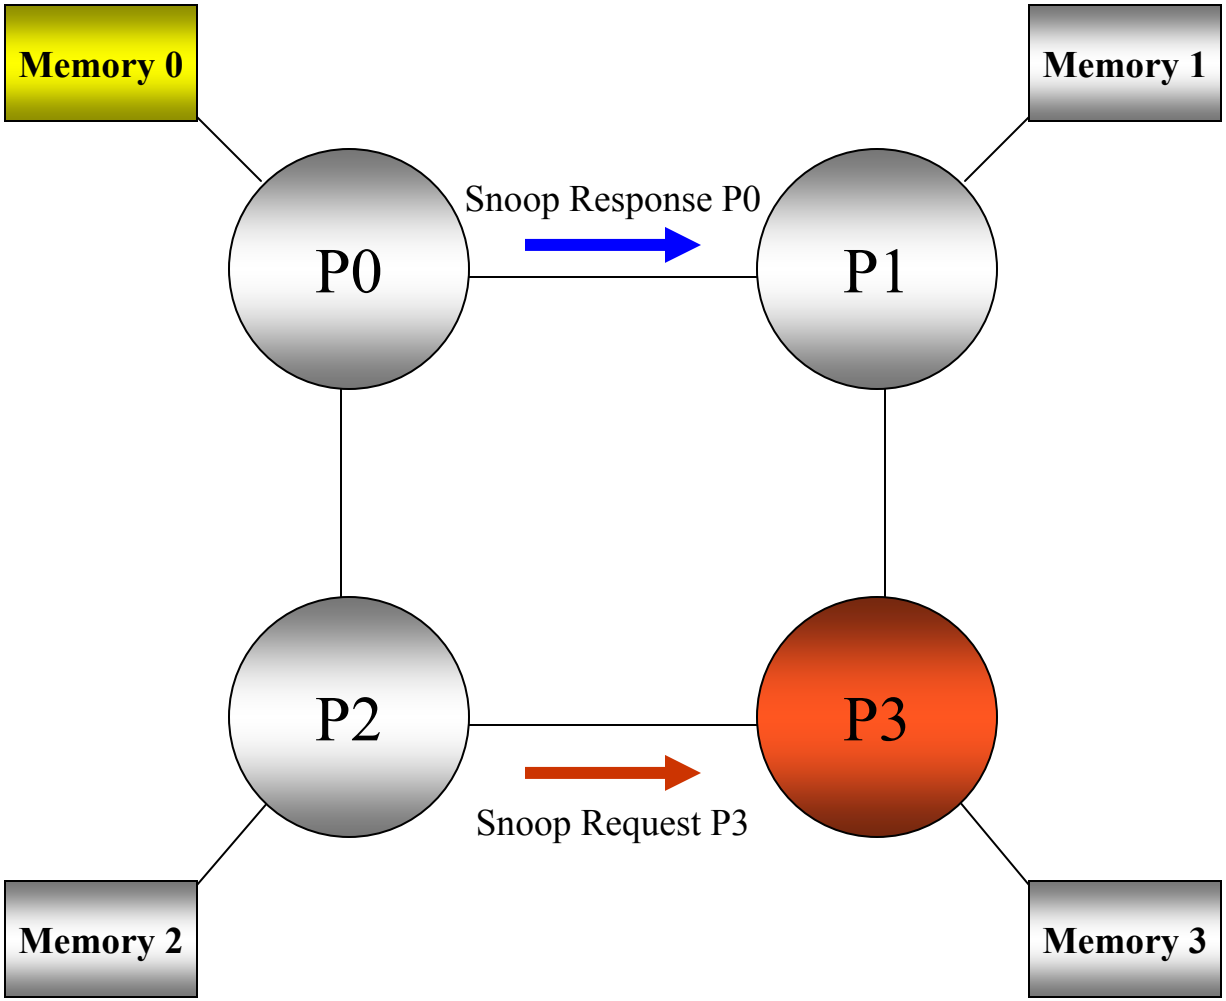

Cache Coherence Protocol Read Transaction Example

Step 5

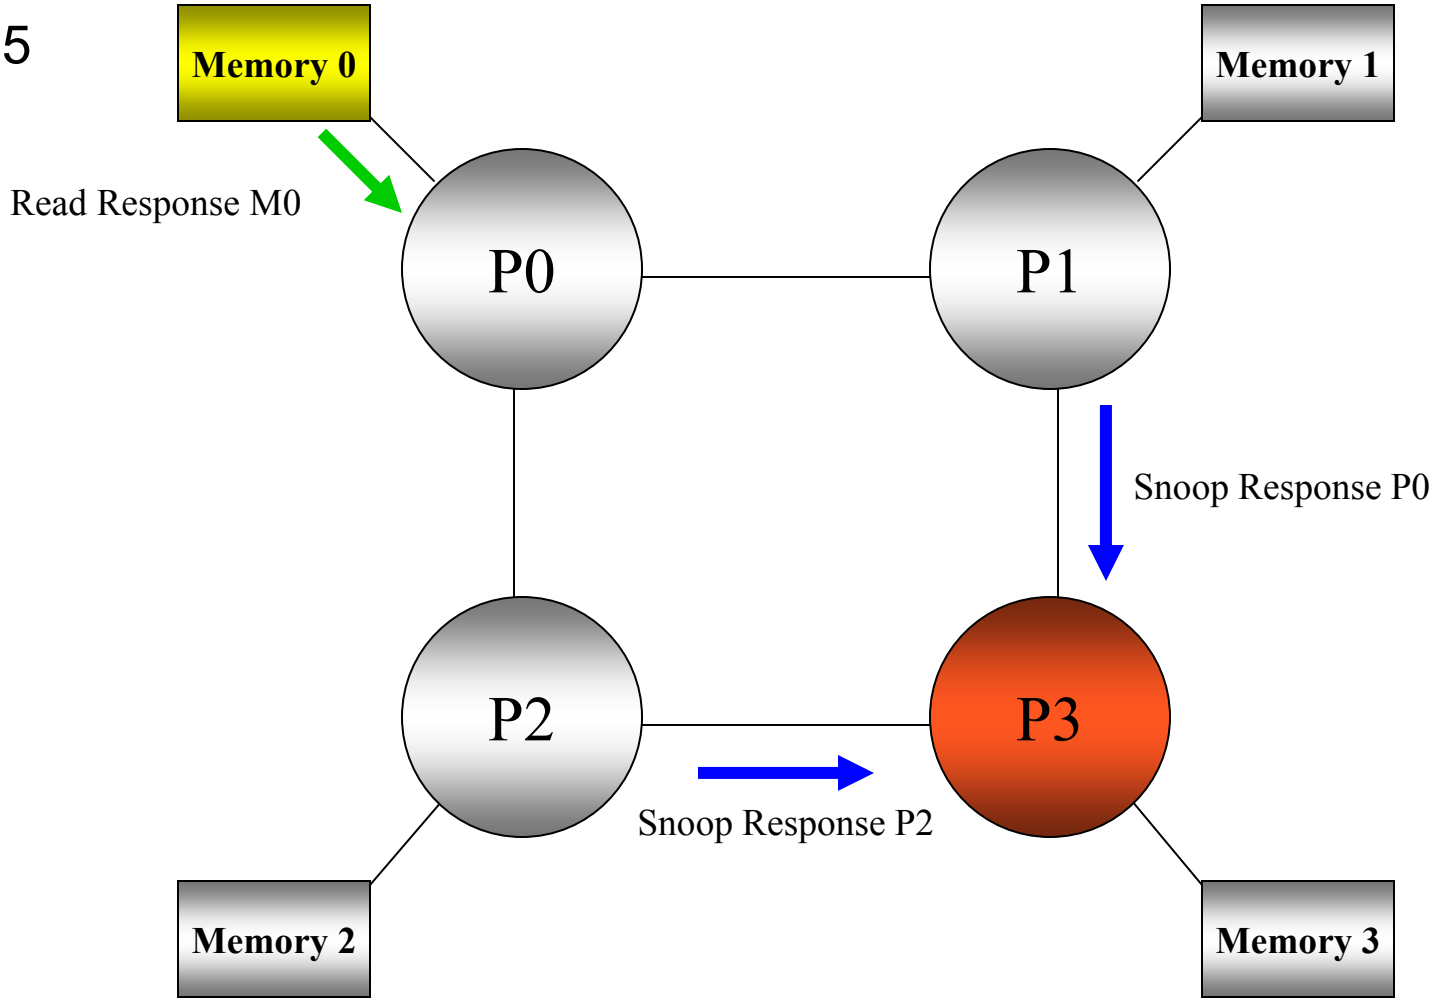

Cache Coherence Protocol Read Transaction Example

Step 6

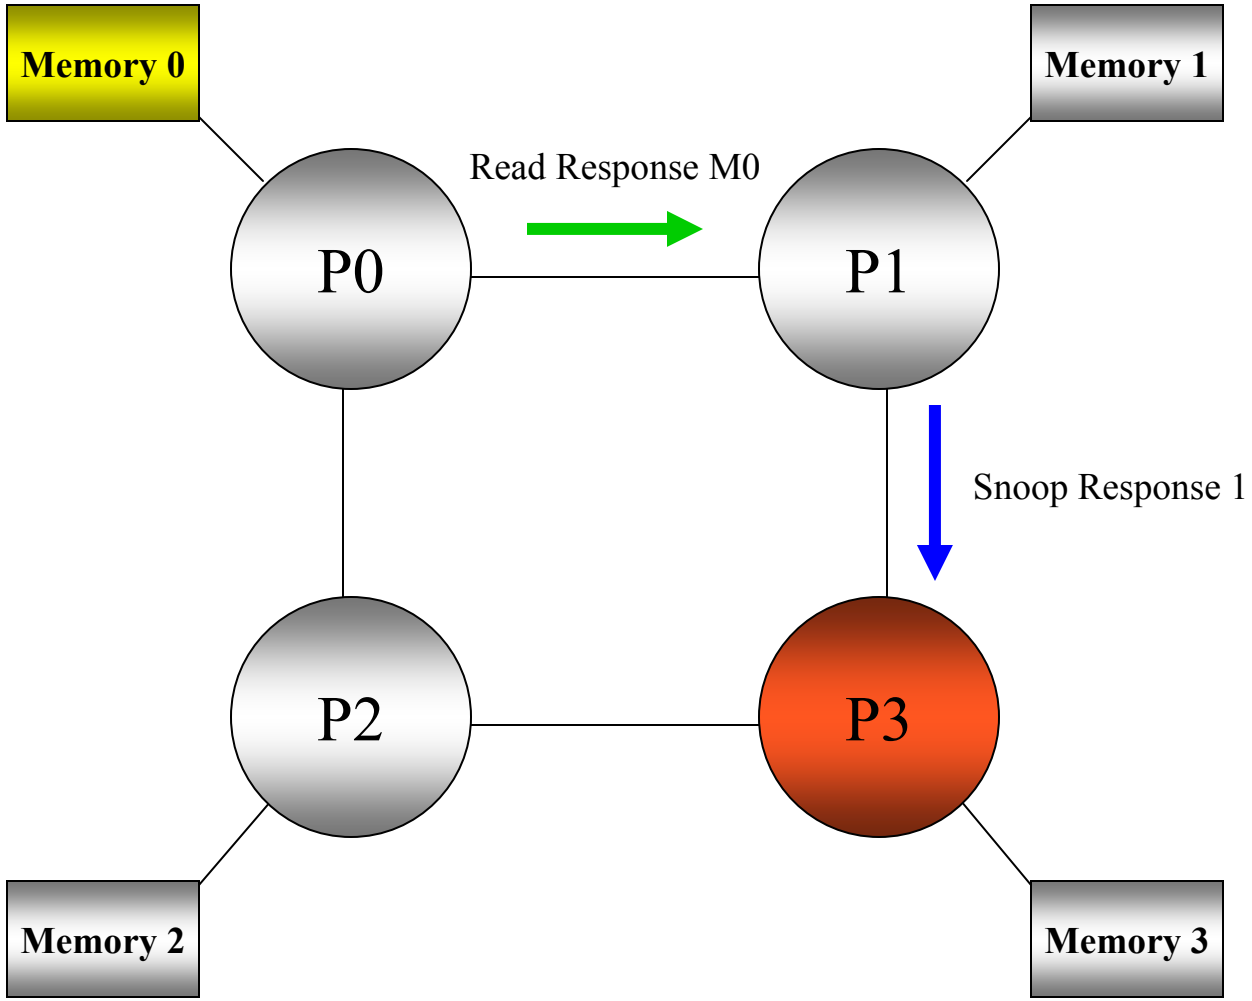

Cache Coherence Protocol Read Transaction Example

Step 7

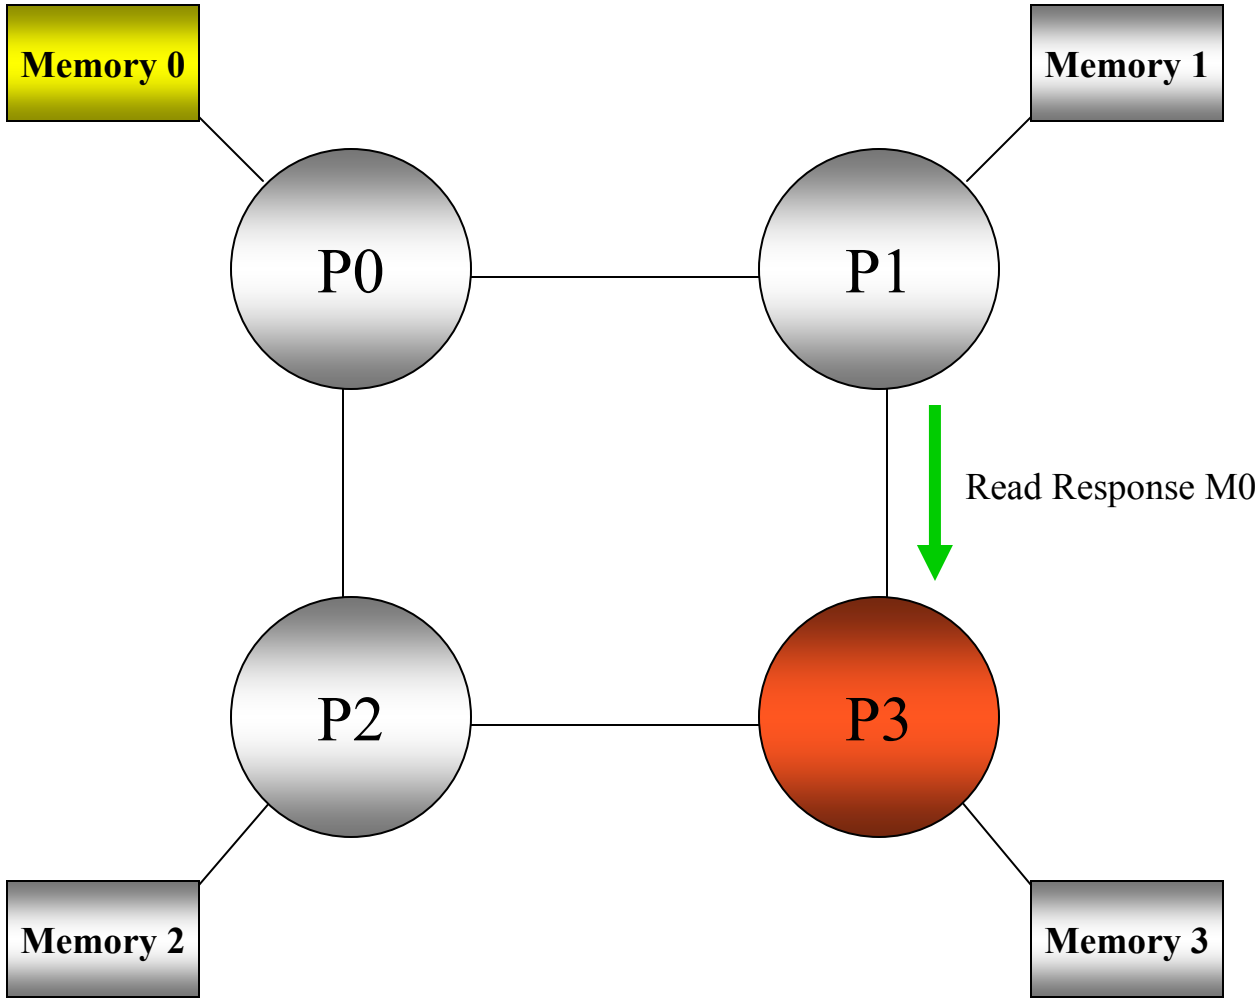

Cache Coherence Protocol Read Transaction Example

Step 8

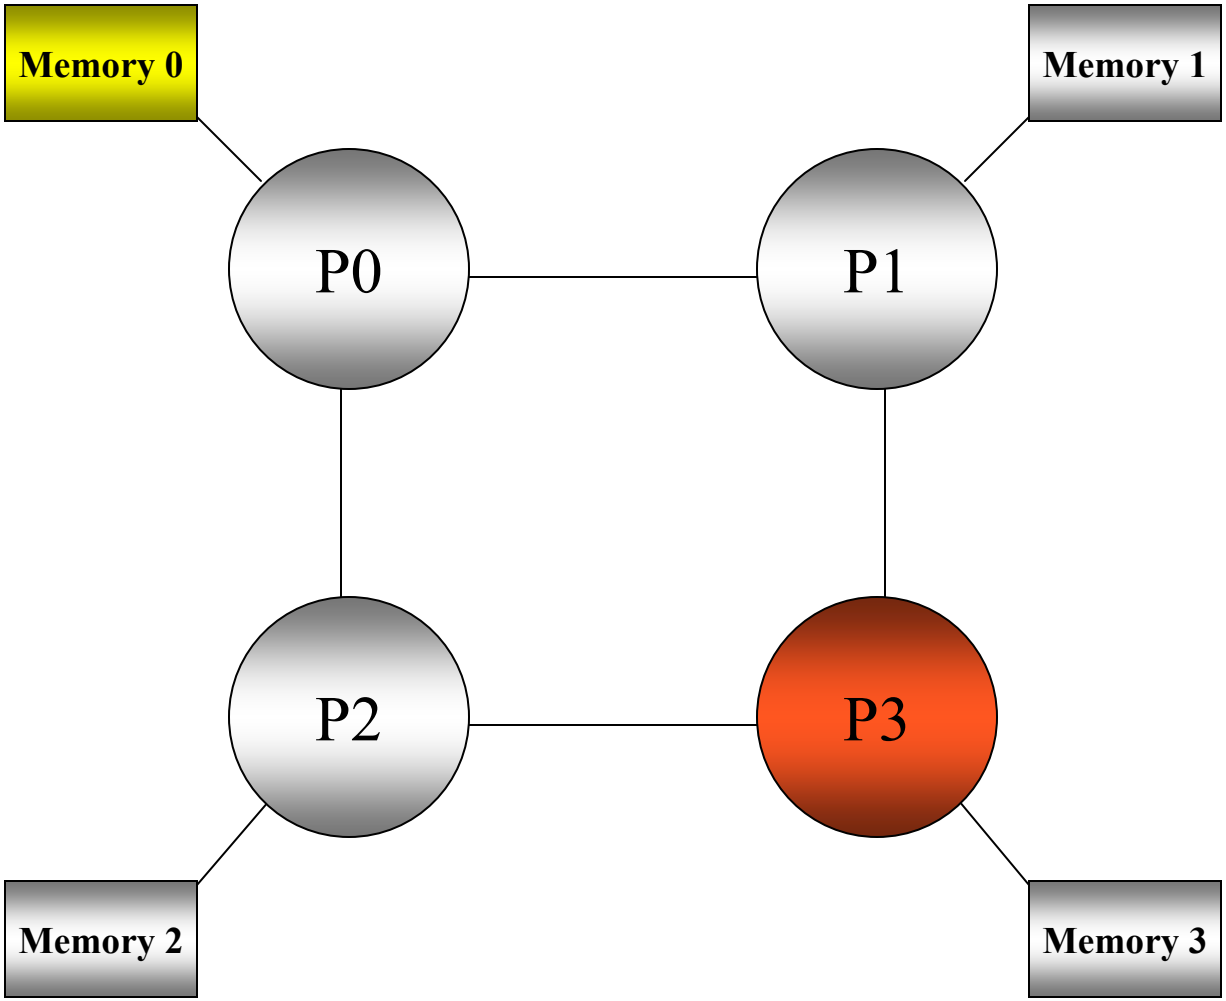

Cache Coherence Protocol Read Transaction Example

Step 9

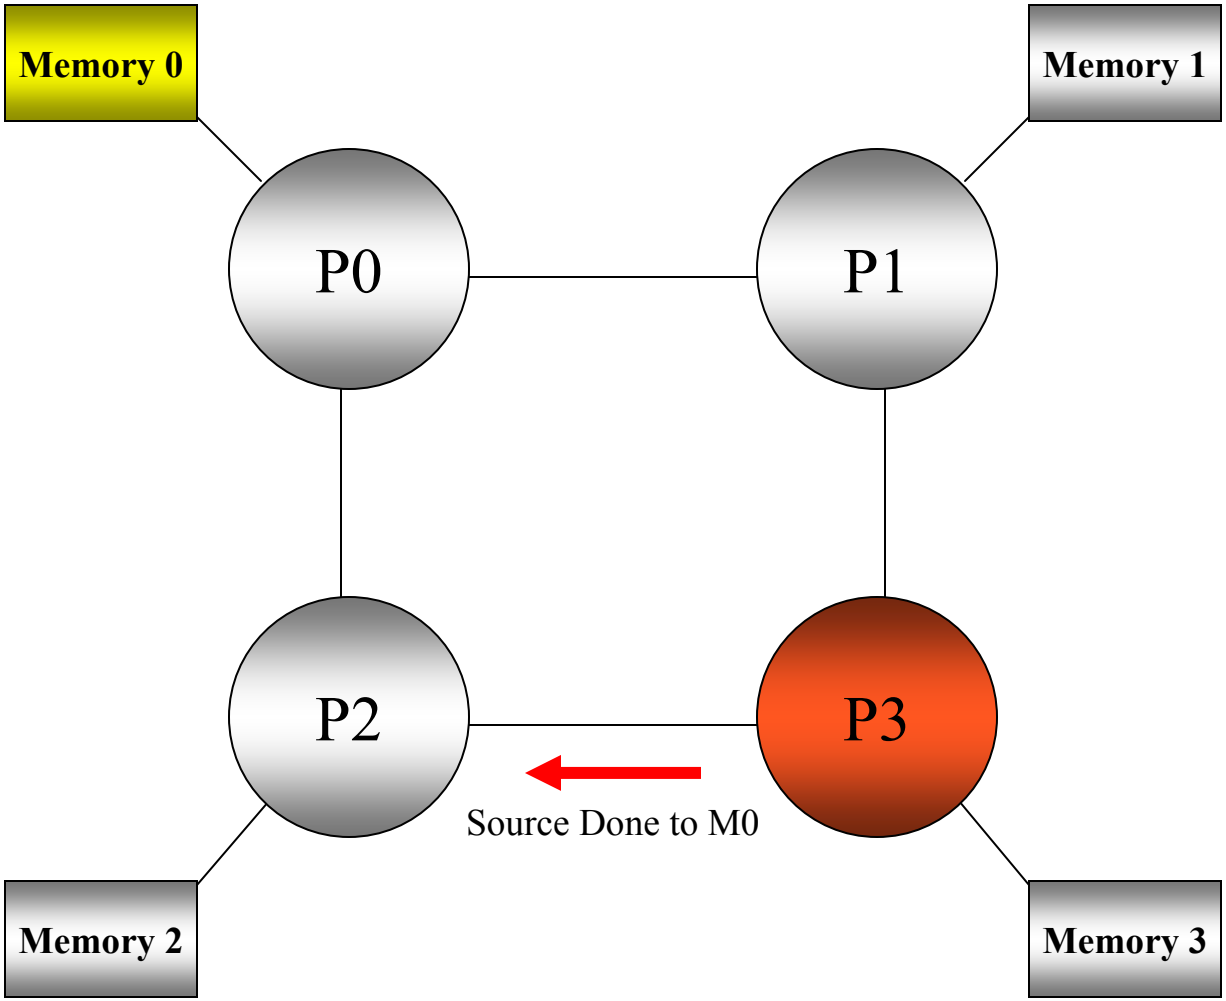

Cache Coherence Protocol Read Transaction Example

Step 10

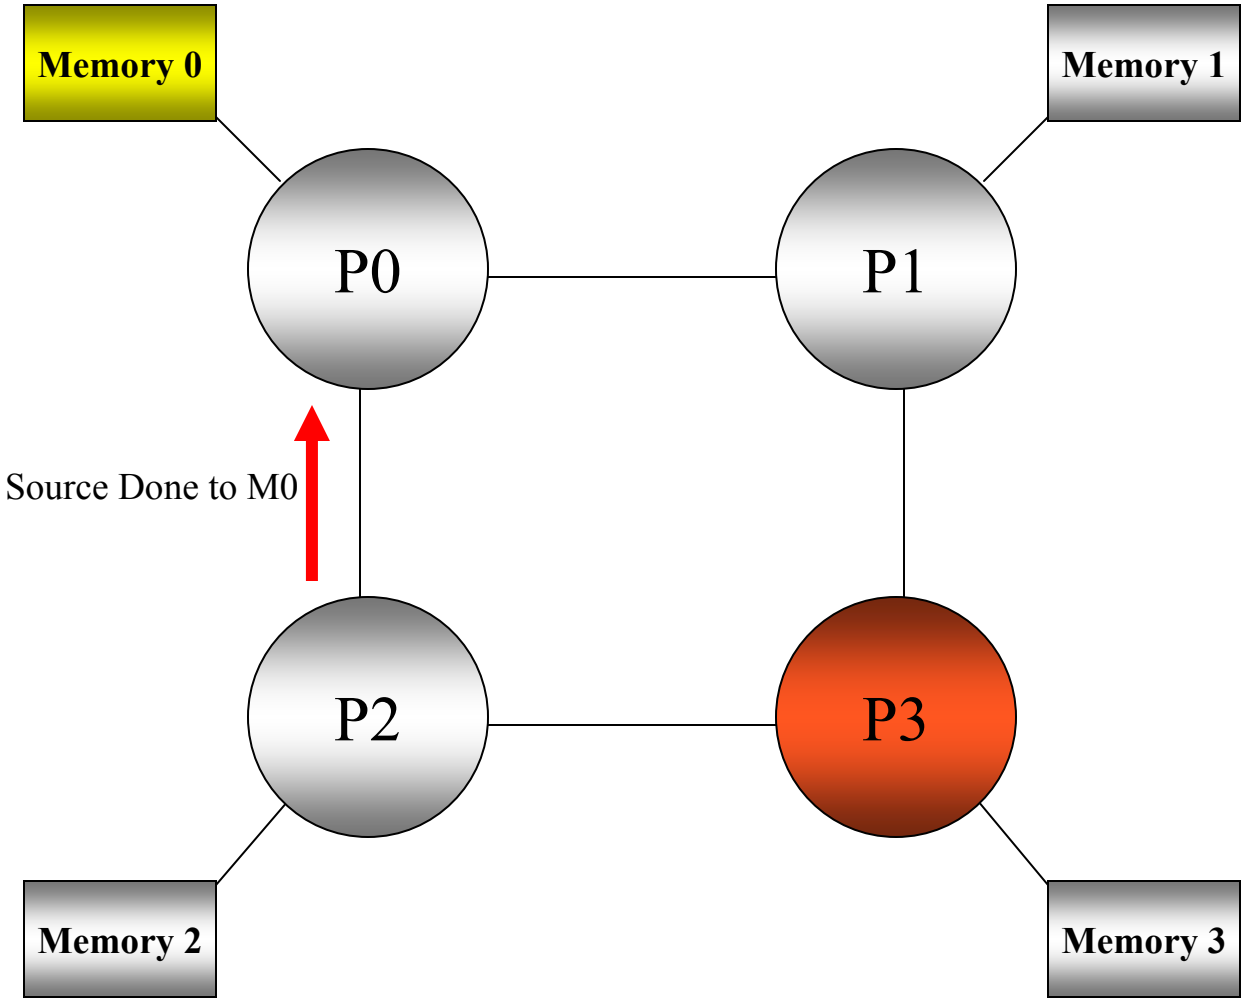

2-way System Topology

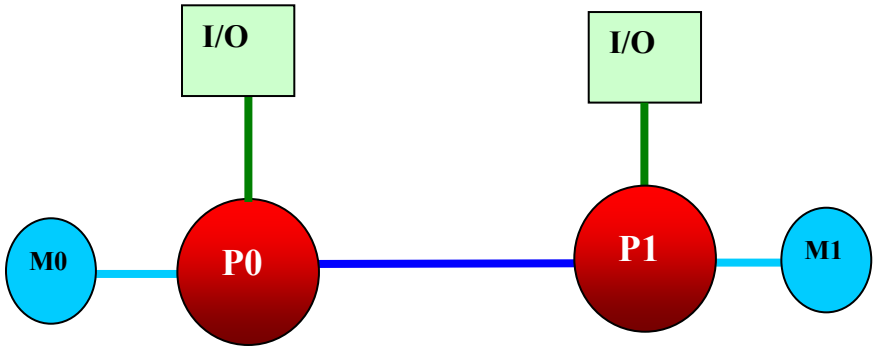

- System parameters
  - 16 DIMMs (up to 32 GB using 256Mb DRAM)
  - 2 HyperTransport™ links available for I/O
  - Bisection-bandwidth = 6.4 GB/s
  - Diameter = 1 hop

4-way System Topology

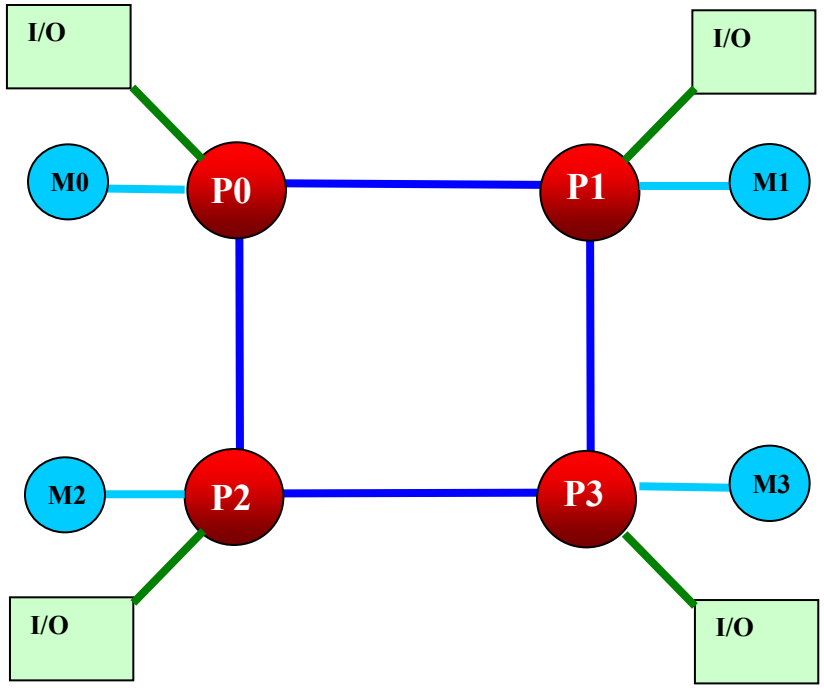

- System parameters
  - 32 DIMMs (up to 64 GB using 256Mb DRAM)
  - 4 HyperTransport™ links available for I/O
  - Bisection-bandwidth = 12.8 GB/s
  - Average-diameter = 1.33 Hops

8-way System Topology

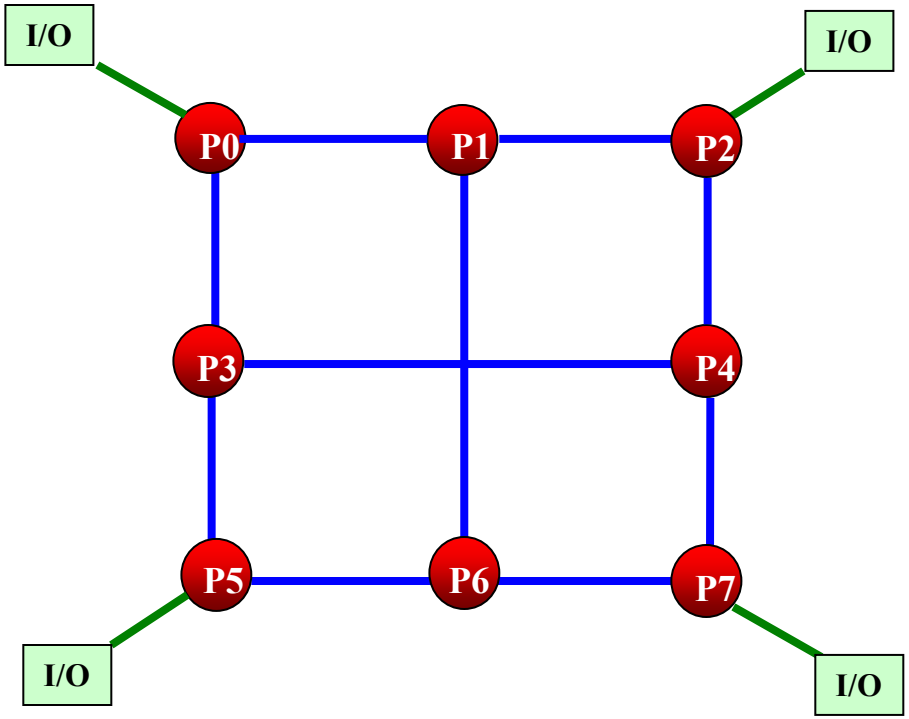

- System parameters
  - 64 DIMMs (up to 128GB using 256Mb DRAM)
  - 4 HyperTransport™ links available for I/O
  - Bisection-bandwidth = 25.6 GB/s
  - Average-diameter = 1.71 hops

Scalability Beyond 8P

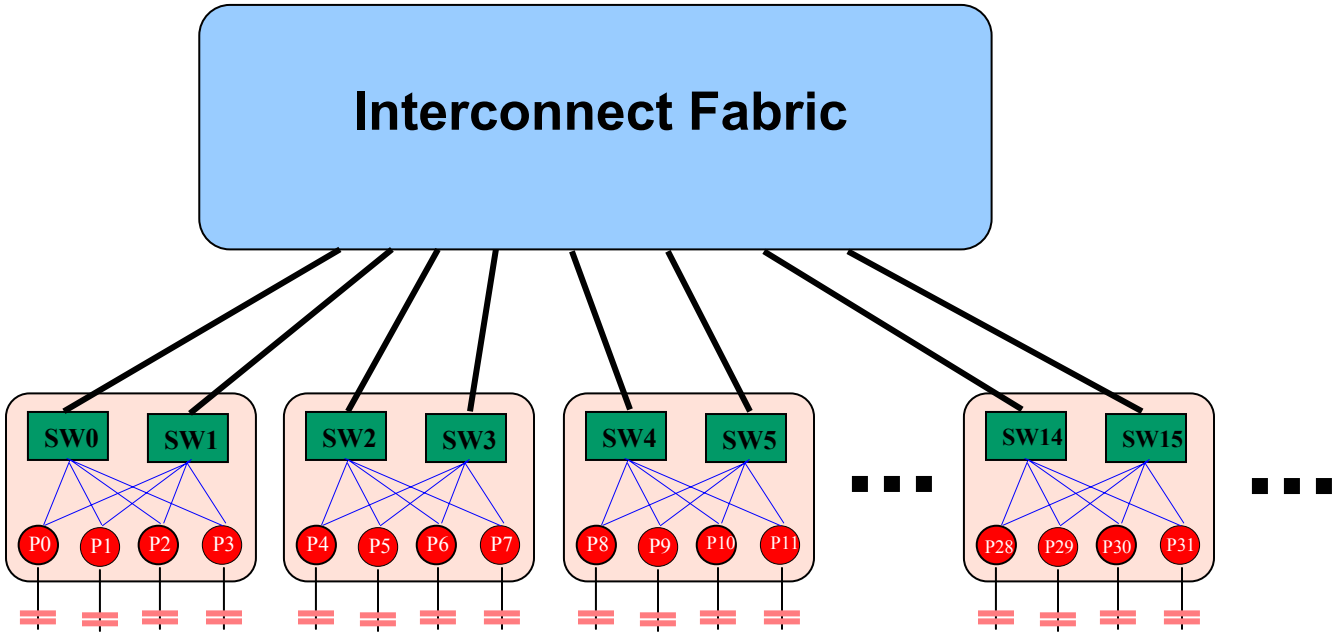

- Scaling beyond 8P is enabled
  - External HyperTransport™ switch
- Coherent Interconnect
  - Snoop filter
  - Data caching

4P Server RDK - Quartet

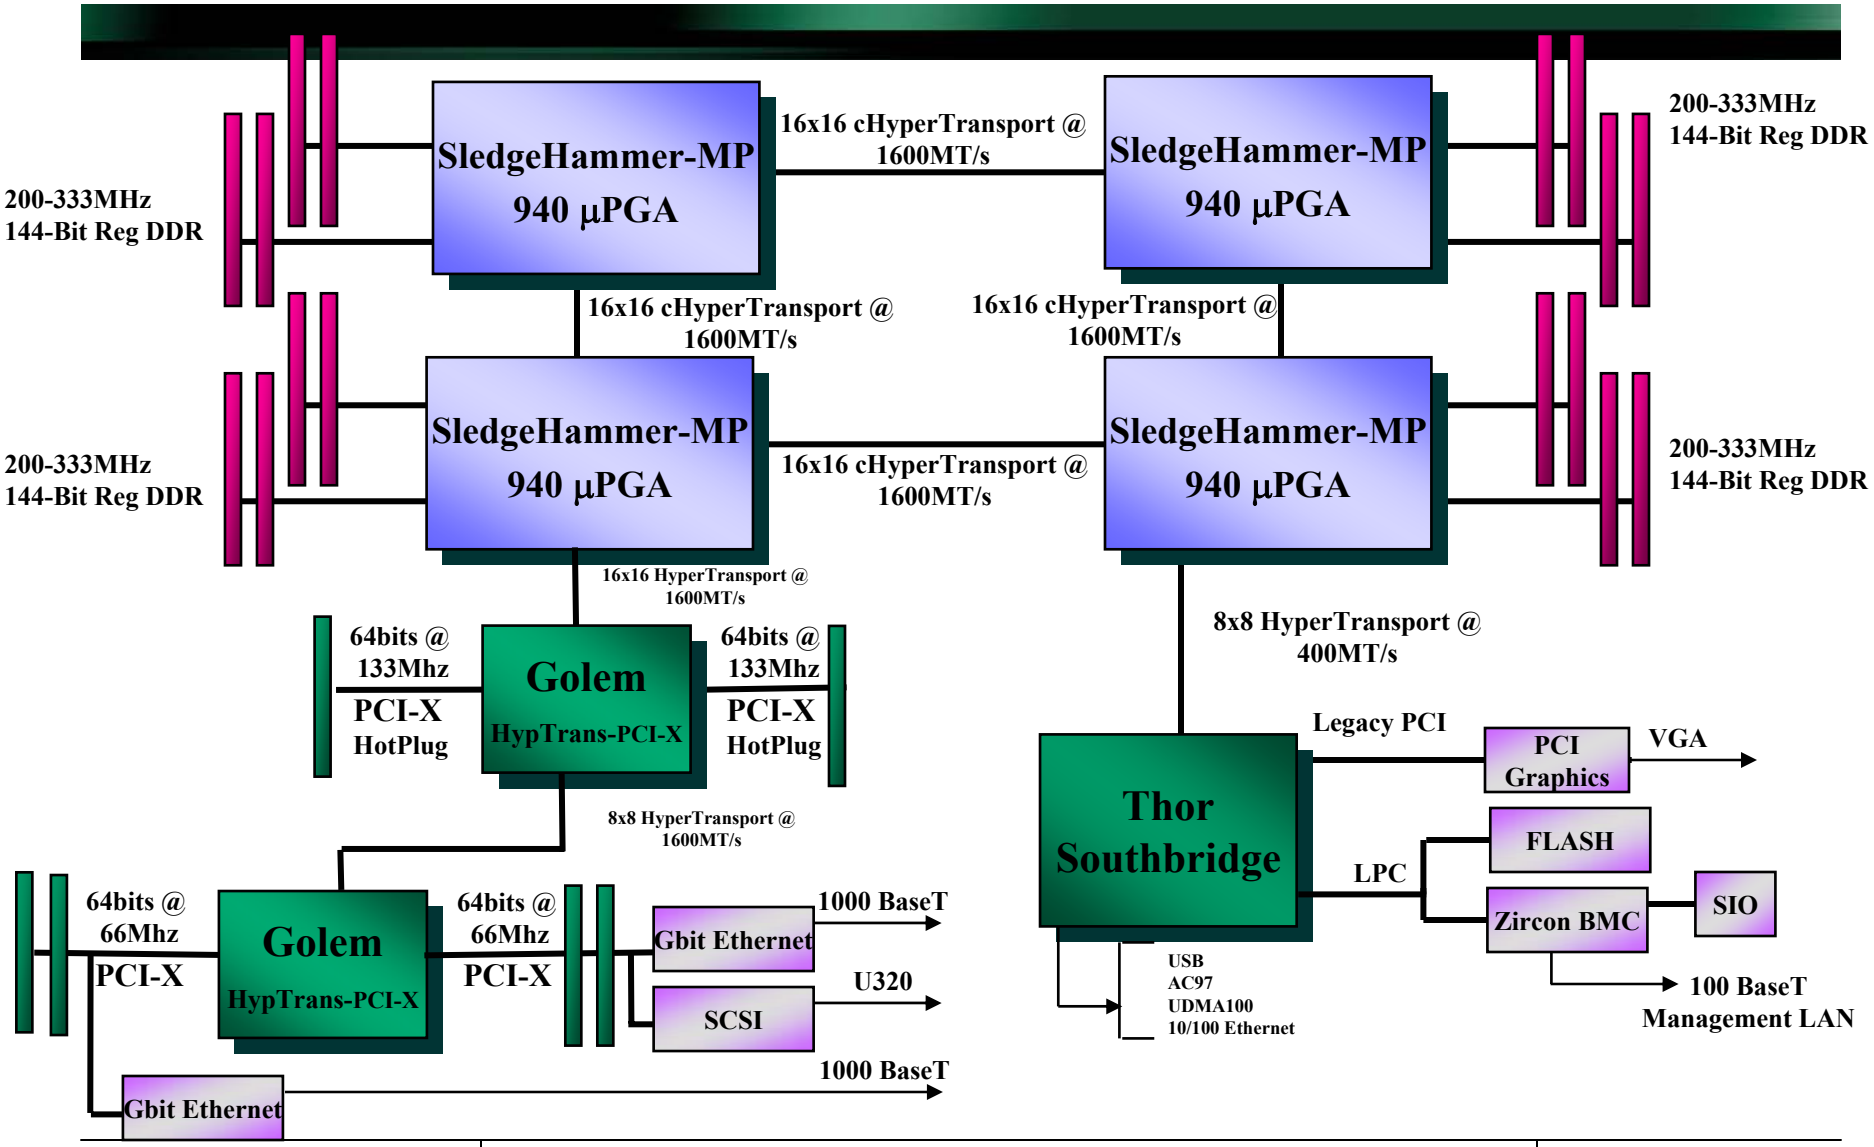

4P System — Board Layout

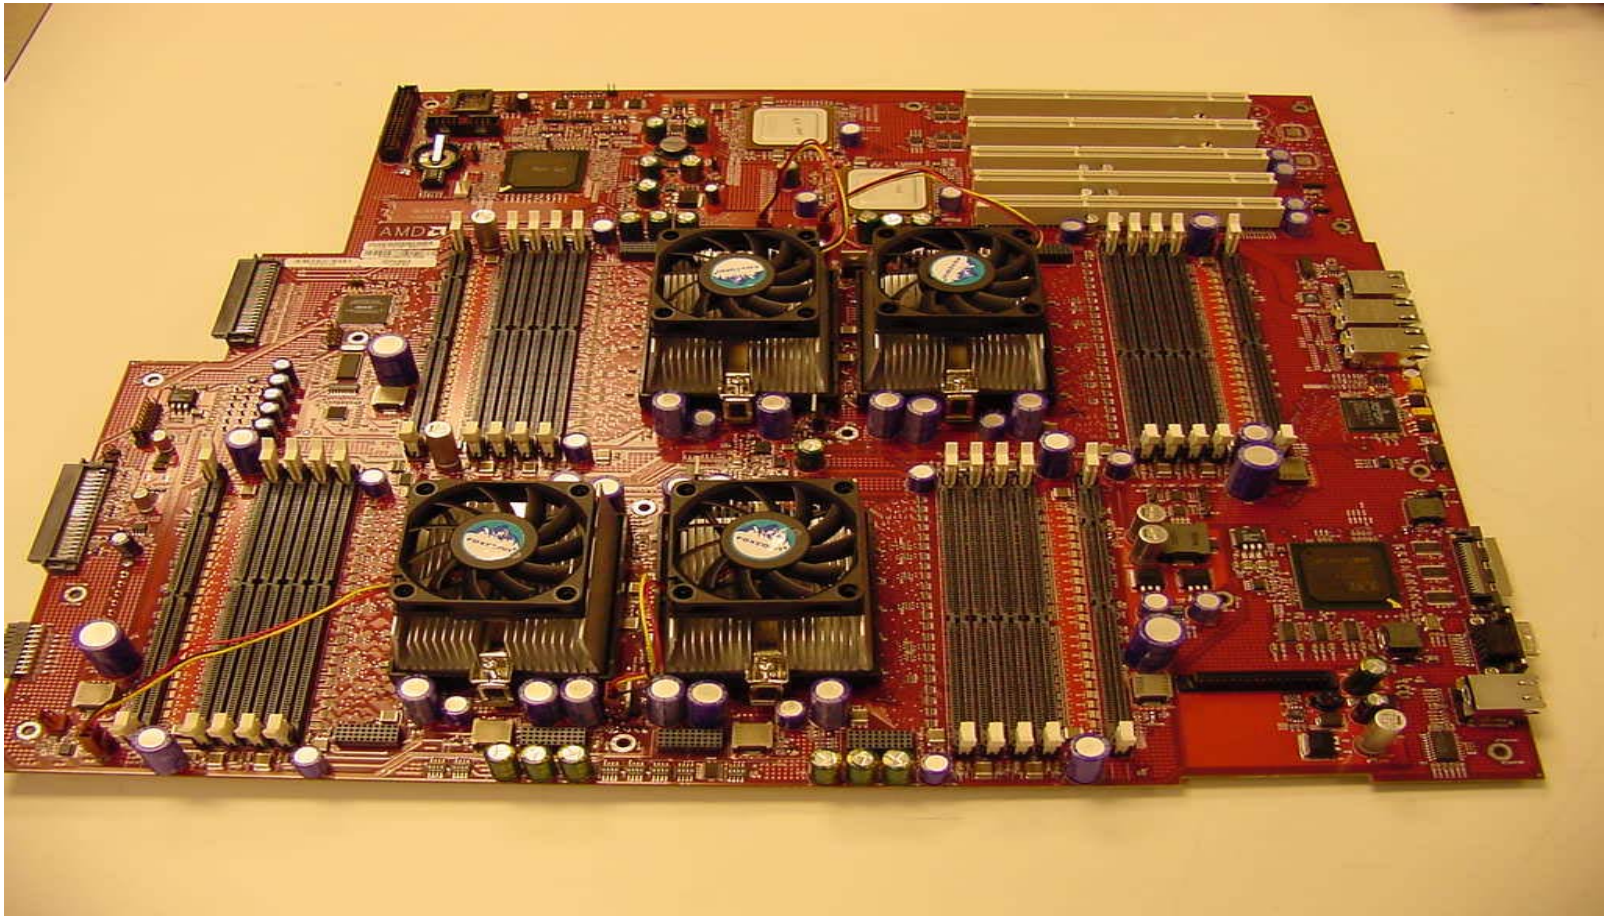

Quartet System Concept (Front-view)

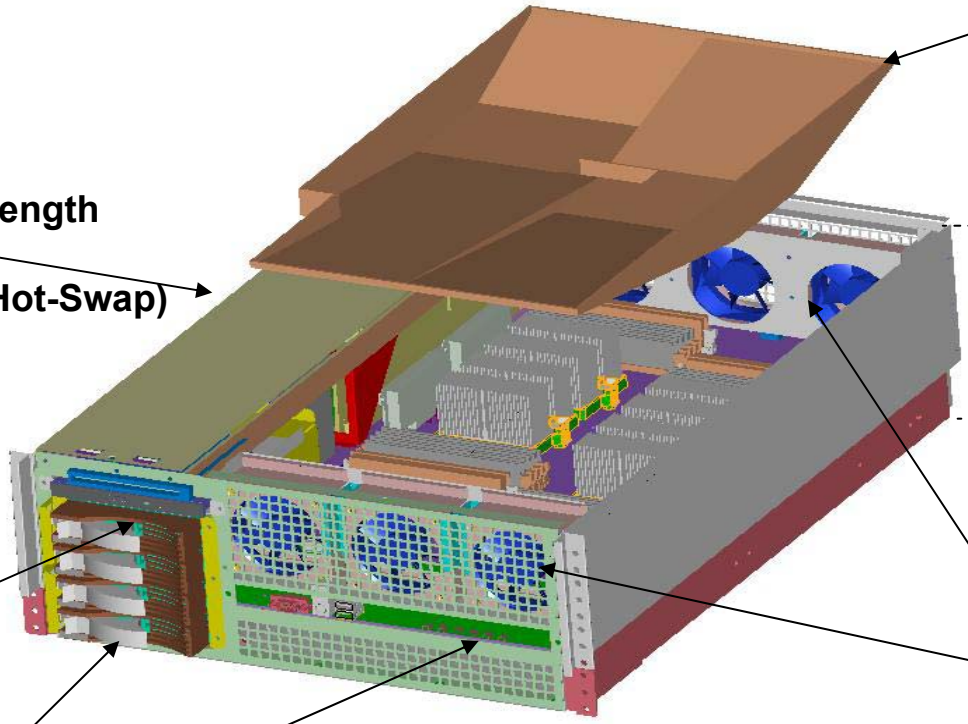

Molded Plenum

6 total Full Length
PCI Cards
(2 ea PCI-X Hot-Swap)
(4ea 64/66)

4U Form factor

Slim line CD &
Floppy

Hard Drive
Bays:
4 eaHotswap 1”
SCSI(OEM Insert)

Front Panel
Personality
Module

6 ea Hot-Swap Fans
Main Planer area
(Remove from top)

Quartet System Concept (Rear-view)

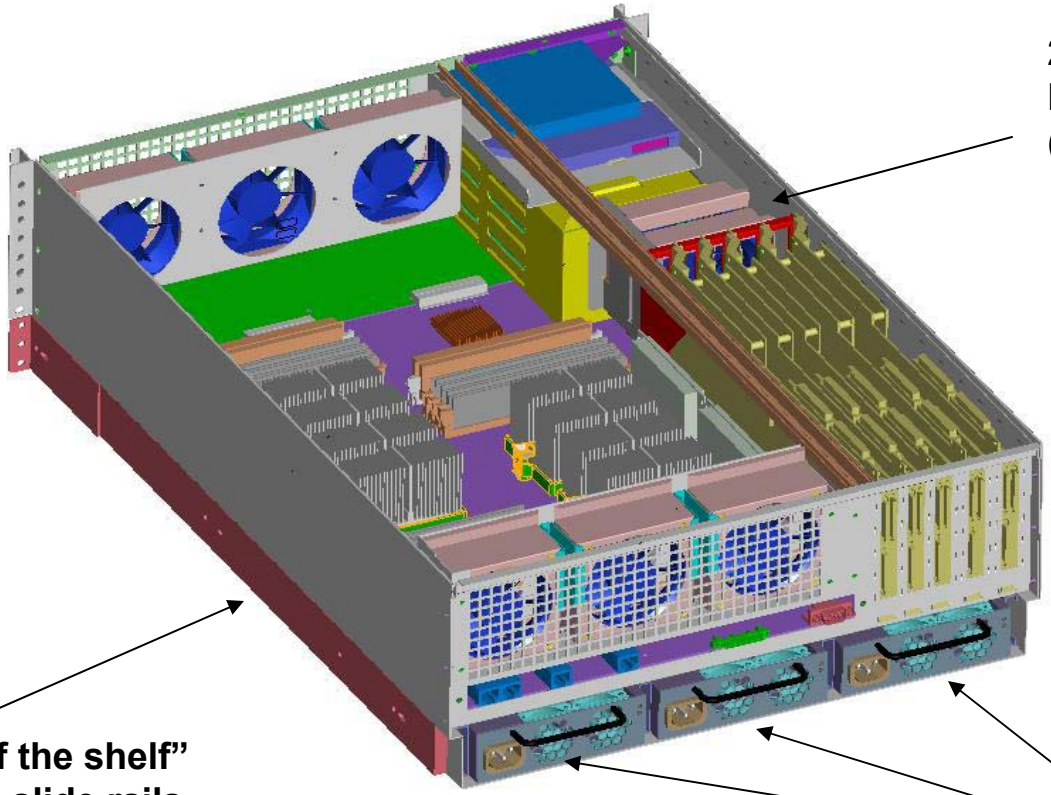

2 ea Hot-Swap Fans
Drive/PCI area
(Remove from top)

Integrated “off the shelf”
Full extension slide rails

Power supplies:
Hot-Swap 2+1 1U
Redundant 500-550W


AMD
Summary

The Hammer Story

AMD DEVELOPER symposium 2002

- The right instruction set
  - Excellent compatibility
  - Excellent performance future
- The right system architecture
  - Great memory and IO capacity and bandwidth
  - Great memory latency
  - Simple “lego” system configuration
- A strong ecosystem
  - Support chips, Software tools, motherboards
- Millions of 64 bit CPUs in 03

Attribution

AMD DEVELOPER symposium 2002